UNIVERSITY of PENNSYLVANIA

PHILADELPHIA 19104

The Moore School of Electrical Engineering D2
DEPARTMENT OF COMPUTER AND INFORMATION SCIENCE
Automatic Program Generation Project

Program Optimization

Based On

A Non-Procedural Specification

By

Kang-Sen Lu

December 1981

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE *(and Subtitle)* Program Optimization Based On a Non-procedural Specification.		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER Moore School Report
7. AUTHOR(s) Kang-Sen Lu		8. CONTRACT OR GRANT NUMBER(s) N00014-76-C-0416
9. PERFORMING ORGANIZATION NAME AND ADDRESS University of Pennsylvania, Moore School of Electrical Engineering, Phiadelphia, PA 19104 Department of Computer Science		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Information Systems Program, Code 437 Arlington, Virginia 22217		12. REPORT DATE December 20, 1981
		13. NUMBER OF PAGES 279 pages
14. MONITORING AGENCY NAME & ADDRESS(If different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) General Distribution	DISTRIBUTION STATEMENT A Approved for public release; Distribution Unlimited	
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Program Optimization, Automatic Program Generation, Very High Level Language Compiler, MODEL		

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This dissertation deals with two related problems: development of a methodology for achieving memory and computation efficiency of computer programs, and the use of this methodology in very high-level programming and associated automatic program generators.

Computer efficiency of programs has many aspects. Usually additional memory saves computation by avoiding the need to recompute certain variables. Our emphasis has been on reducing memory use by variables sharing

memory space, without requiring recomputation. It will be shown that this also reduces computation overhead. The most significant savings are due to sharing memory in iterative steps. This is the focus of the reported research.

The evaluation of memory use of the many possible alternatives for realizing a computation is highly complex and requires lengthy and expensive computations. We have developed a heuristic approach, which has been very effective in our experience, and which is practical and economical in use of the computer. Basically it consists of evaluating global memory usage alternatives on each level of nested iteration loops, starting with the outside level and moving inwardly. Thus we neglect the rare impact of a nested iteration loop on the memory usage calculated for an outside iteration. This has lead to the principle of maximizing size of loop scopes in a program as a means to attaining a more efficient program for present-day sequential computers.

The automatic design of efficient programs is also essential in use of very high level languages. The use of very high level languages offers many benefits, such as less program coding, less required proficiency in programming and analysis, and ease in understanding maintenance and up-dating of programs. All these benefits are conditioned on whether the language processor can produce satisfactorily efficient program.

The dissertation reports the design and implementation of a new version of the MODEL language and processor which incorporates algorithms for producing more efficient programs. The dissertation describes briefly the MODEL non-procedural language and the analysis, scheduling, and code generation tasks.

PROGRAM OPTIMIZATION
BASED ON
A NON-PROCEDURAL SPECIFICATION

Kang-Sen Lu

A DISSERTATION

in

Computer and Information Science

Presented to the Graduate Faculties of the University of Pennsylvania
in Partial Fulfillment of the Requirements for the Degree of Doctor of
Philosophy.

1981

Supervisor of Dissertation

Graduate Group Chairperson

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to my advisor, Professor Noah Prywes, for his support and patient guidance during the period of this research. His advice and encouragement have made completion of this dissertation possible.

Professor Amir Pnueli deserves special thanks for the many valuable discussions during his intermittant visiting research in the University of Pennsylvania.

I also would like to thank Maya Gokhale and Wu-Hung Liu for their dedicated reading of the drafts, and suggestiions for improvements.

Finally, I wish to express my gratitude to my wife, Yu-Chen, for her encouragement, support, and understanding throughout the entire research.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

1.1 OBJECTIVES OF THE RESEARCH

This dissertation deals with two related problems: a) development of a methodology for achieving memory and computation efficiency of computer programs, and b) the use of this methodology in Very High-Level programming Languages (VHLL) and associated automatic program generators.

There are many aspects to computer efficiency of programs and we had to be selective in choosing to focus our research on the aspects that we considered most important. Optimization of computer efficiency of programs concerns the two major aspects of reducing computation time and reducing usage of memory space. We have selected the memory space reduction aspect for two reasons. First, the excessive use of memory has been the major disadvantage in use of VHLLs, especially where interpreting techniques have been used in the language processor. Second, as will be shown, reduction

of memory space also reduces computation overhead. Further we have not considered techniques which save memory through recomputing of some variables as the impact of such techniques on computing time may be enormous. The potential for reducing use of memory exists through both global and local analysis of a program. Among the many methods for reducing memory use, we have emphasized global methods for reducing memory use particularly through sharing memory space by variables in iterative steps of the program. This approach represents the potential for the most siginificant savings in memory. In summary, the dissertation concerns reduction in use of memory in performing computations specified in a VHLL, particularly through sharing of memory in program iterations.

In most VHLL systems, memory use is determined primarily on a dynamic basis at run time. This is particularly typical of interpreters for VHLLs. The dissertation will show that a global analysis of the VHLL can lead to prescheduling the use of memory and compiling a program which uses memory efficiently. The use of this method can eliminate the most important drawback on use of VHLLs, i.e., the inefficiency in performing the computation.

The evaluation of the many possible global and local alternatives of memory use for realizing a computation is highly complex and requires lengthy and expensive computations. We have developed a heuristic approach, which

has been very effective in our experience, and which is practical and economical in use of the computer. We have generally used the principle of maximizing size of loop scopes in a program as a means for attaining a more efficient program for present day sequential computers. Further, program design decisions are based on evaluation of memory usage alternatives on each global level of nested iteration loops in a program, starting with the outside level and moving inwardly. Thus we neglect the rare impact where memory usage in a local nested iteration loop requires reversing the more global design of the outside iteration loop.

In a VHLL the user can specify the computation more abstractly, i.e. without concern for the efficiency of the algorithm for performing the computation. This contrasts with programs written in lower level languages. Therefore starting with the higher level specification allows the global optimization of the program.

The MODEL VHLL and processor have been chosen in this dissertation to study the optimization problems. The MODEL language is non-procedural. It includes the use of arrays and records data structures which are used widely in both mathematical systems and in data processing. Yet the language is simple enough.

The result of the research has been the incorporation of novel optimization techniques in the MODEL automatic program generator. The new system automatically designs and generates high level language programs, in PL/I, with efficient loop control and economical memory usage, without the user's concern for efficiency of memory allocation. The resulting system demonstrates that an efficient implementation of computations based on a very high-level non-procedural specification is possible and therefore that the use of VHLL can be made practical.

Apart of the questions of incorporating efficiency while generating a program automatically based on a VHLL specification, there are the more basic methods of analysis for improving efficiency of programs. These have been the other objective of this research, i.e. to develop analytical methods for determining how a conventional program can be made more efficient and to offer methods to determine program design decisions.

1.2 CONTRIBUTIONS

This dissertation addresses the problem of generating efficient programs based on a very high-level non-procedural specifications of the programs. The program optimization uses appropriate algorithms for implementing required computations. Program loop optimization and memory

optimization are the major concern of the research.

More specific achievements include the following results:

1. Methods for semantics analysis of a program specification to develop the information needed for program generation. This includes precedence relationships among program events and indicated order of nesting of loops.

2. Criteria for including events or computations in loops of programs. The approach is to maximize scope of loops as means for reducing memory use and computation time. Repeating program events or computations which satify the following conditions may be included in the scope of a loop: a) the same or related range of iterations, b) continuity of dependencies among the events in the scope of a loop, c) compatibility of a "distinguished dimension" in the many dimensions of repeating events, and d) a conditioned block of events of related ranges can be placed within a loop to further extend the loop scope.

3. A method for determining whether memory space for an array dimension has to be physical or virtual, i.e. whether memory can be shared.

4. A method for evaluating "memory penalty" of selected loop scopes as a basis for choice of the most economic loop design.

1.3 ORGANIZATION OF THE DISSERTATION

The dissertation is divided into seven chapters. The introduction is given in this chapter. Chapter 2 surveys related research, in the fields of programming languages, automatic programming, and program optimization. Chapter 2 is divided into respective sections which deal with procedural High-Level Languages (HLLs), VHLLs, and program synthesis, including their efficiency considerations. The reading of this chapter may be omitted by reader familiar with the state of the art in programming.

Chapter 3 describes the syntax and semantics of the MODEL language. Since its denotational semantics can be found in [SANG 80], the description is from the user's point of view and this chapter can be used as a user's guide.

Chapter 4 describes the semantic analysis done by the MODEL processor. This includes checking for various aspects of inconsistency and incompleteness of the program specification, and correcting the tolerable incompleteness. Most importantly, this chapter describes the internal representation of the programspecification, including discovering the precedence relationships among the program entities, by an Array Graph.

Chapter 5 discusses the range propagation method which classifies all the array dimensions and assertion subscripts into range sets according to their respective ranges (i.e.

number of repetitions) and corrects omission of subscripts. The range sets will be the candidates for loop construction.

Chapter 6 discusses the major contribution of the research, the scheduling algorithm, whose function is to synthesize a computation procedure. The algorithm generates design of an optimized program. The program optimization is achieved by maximizing the loop scopes, selecting loops of the least memory use, and merging the loops of related ranges.

Chapter 7 discusses the code generation. Code generation is a process which takes the program schedule as input and generates a PL/I program ready for compilation.

Suggested future work is presented in Chapter 8.

The detailed documentation of the system is rather lengthy and has not been included in this dissertation. A report documenting the entire MODEL system has been prepared by the author separately from the dissertation. Also program listings further document the research. The system has been subject to extensive experimentation and examples of specifications and resulting automatically generated programs are given in the appendix.

CHAPTER 2

SURVEY OF RELATED WORK

It has been stated that "almost anything in computer
science can be made relevant·to the problem of helping to
automate programming"[FELD 72]. Therefore any survey of
programming language development must be in some respect
incomplete. An excellent overall discussion of the trends
in software development research can be found in [WEGN 79].
The survey of the recent research in this chapter emphasizes
the fields of programming languages, automatic programming,
and program optimization, which are the major interests of
this thesis. The survey includes a review of the impact of
problems of efficiency on programming and the relevance of
the reported research to these problems.

Among the approaches suggested to date to improve the
quality of the software development are: modularity, strict
type checking, data abstraction, higher level operations and
general data structures, non-procedurality, and domain
specific languages. Each of these has been successful in

some aspects. In the following we classify programming languages and systems into three categories, namely procedural high-level languages, very high-level languages, and automatic program synthesizing systems. From each category a few representative languages which incorporate some of these concepts will be briefly reviewed.

2.1 PROCEDURAL HIGH-LEVEL LANGUAGES

Procedural high-level languages provide control statements for the user to compose efficient programs. The user specifies the computation in a procedural way, which is usually tedious and prone to error. The need for a flowchart to help the programmer analyze and document the program logic shows that procedural programming could easily confuse even the program designer. The structured programming discipline has been advocated in writing programs, and linguistic features such as type checking and abstraction mechanisms were suggested to further reduce errors by programmers.

2.1.1 EXAMPLES OF HIGH-LEVEL LANGUAGES

The programming language PASCAL and its derivatives are examples of procedural HLLs. They emphasize type checking at compilation time to catch erroneous uses of data as early

as possible. The type of an object is characterized by the set of values that the object can assume, and the set of operations that may be performed on the object. Primitive data types are predefined in the programming language. Users may define new data types from primitive data types or from other user-defined data types. Since it is required to associate types with variables and parameters of subprograms, objects with distinct properties are clearly distinguished in a program by their data types and the distinction is enforced by the compiler. It has been claimed that requiring typed objects contributes to program reliability. Many programming languages have followed the spirit of PASCAL in strict type checking. For example, MESA[GEMS 77], and ADA [ADAA 79] are typed languages. Although type checking is claimed to be a powerful tool for increasing software reliability, it is realized that the benefit from the linguistic mechanisms do not come automatically. A programmer must learn to use them effectively. Also it is not always desirable to remain within the type checking system because sometimes the violation is logically necessary, especially in the area of systems programming. For example, a compile-and-go system will have to convert the type of a generated object code from <u>data</u> into <u>procedure</u>. The answer has been to make those occasional type violations as explicit as possible. Therefore, these type violations are less dangerous since they are clearer to the reader.

Abstraction has long been suggested as helpful in programming methodology. Many conventional languages have supported procedural abstraction with functions and subroutines. The class concept of SIMULA has pioneered in data abstraction. Parnas[PARN 72] also pointed out that the criteria of decomposing a software system should not be based on the steps of the algorithm, but instead, a module in a decomposed system should be characterized by its knowledge of some design decisions which it hides from others. Its interface or definition should be chosen to reveal as little as possible about its inner workings.

The programming language CLU[LSAS 77] was designed to support the use of abstractions in program construction. In CLU, each object has a particular type. A type defines a set of operations that create and manipulate objects of that type. The basic data abstraction mechanism of CLU is the <u>cluster</u> which is used to define abstract data types. The cluster provides a representation for objects of certain type and an implementation for each of the operations. The type checking done for assignments and argument passing ensures that the behavior of an object is indeed characterized completely by the operations of its type.

The language ADA[ADAB 79] has been designed with the concern of program reliability and maintenance. Program variables are required to be declared with their types. Automatic type conversion is prohibited. Thus, compilers

can ensure that the types of objects satisfy their intended use. Modules in ADA allow the specification of groups of logically related entities. In their simplest form modules can represent pools of common data and type declarations. In addition, modules can be used to describe groups of related subprograms and encapsulated data types, whose inner workings may be concealed and protected from their uses. A module is generally provided in two parts: a module specification and a module body with the same identifier. A module specification may contain the specification of subprograms which are visible to the other program units. The implementation of the subprograms is declared in the module body, and it is not accessible outside the module. As a consequence, a module with a module body can be used for the construction of a group of related subprograms, where the logical operations accessible to the user are clearly isolated from the internal entities.

Because of the distinction between abstractions and implementations, data abstractions ease program modification, maintenance, understanding, and verification. However, the quality of any program depends upon the skill of the designer. In a programming language supporting data abstraction the skill is reflected in the choice of abstractions. Abstractions should be used to simplify the connections between modules and to encapsulate decisions that are likely to change.

2.1.2 COMPILER OPTIMIZATION

The concern over the inefficiency of compiler generated code dates back to the early introduction of high-level programming languages. Program optimization techniques have been incorporated into compilers to produce more efficient code. The efficiency of a program may be measured using various aspects, such as the execution time of the code, the size of the code, or the size of the data area. The emphasis in program optimization may depend on the characteristics of respective computer architecture or programming language.

Optimization techniques for high-level languages such as FORTRAN or PL/I emphasize code optimization, i.e. producing better object code than the most obvious one for a given source program. The efficient utilization of the registers and instruction set of a machine can improve program efficiency significantly. Most issues in this area are highly machine dependent. Optimization techniques which are not machine dependent include identifying common subexpressions and moving loop invariant computation outside of the loop.

Code optimization techniques are generally applied before or during the code generation phase of a compilation process of a HLL program. The major issues in the code generation phase are deciding what instructions to use, in

what order to execute, and where to store the intermediate results in temporary storages. Bruno and Sethi[BRSE 76] showed that the problem of generating minimal length code for a one-accumulator machine is NP-complete problem. However, if there are no identified common subexpressions in an arithmetic expression, it is possible to generate optimal code in linear time[AHJO 76]. In the presence of common subexpressions, some heuristic algorithms may be used to produce code that in the worst case is three times as long as optimal[AHJO 77].

Many optimization techniques have been found to be machine independent. These include constant subsumption, common subexpression suppression, code hoisting, and dead code elimination. These techniques usually need information that can only be obtained by a global analysis of the program. The global flow analysis finds the related definitions for a use of a variable and the related uses for a definition of a variable. A formal discussion of the global analysis can be found in [SCHA 73]. [AHUL 78] contains a rather complete survey of code optimization techniques.

Recent research interest in compiler design has shifted to the automation of the code generation phase. A table-driven approach has been proposed by Susan Graham[GRAH 80]. The description of machine instructions is encoded in a table used by the code generator where the

function of each instruction is represented by a tree. The input to the generator is a subprogram in a tree representation. When a subtree in the program matches some instruction tree, the corresponding instructions are emitted. Thus, the task of code selection is reduced to a symbolic pattern matching problem. The advantages of this approach include the ease in modifying the code generator for a new machine and thorough search of the instruction set even if the target machine has an asymmetrical instruction set.

The Production-Quality-Compiler-Compiler (PQCC) project at Carnegie-Mellon University has aimed at building a truly automatic compiler writing system[LCHN 80]. The system generates a compiler from descriptions of both the source language and the target computer. The emphasis of the investigation is on the code generation phase. In order to keep the PQCC system general only the optimization techniques which can be parameterized for different machine architectures are included in the system. The machine dependent optimizations are isolated in such a way that only the tables may contain machine dependent information but the procedure code which operates with the tables is machine independent. The objective of the project has been to obtain simultaneously the retargetability and a high level of optimization of a compiler.

2.2 VERY HIGH-LEVEL LANGUAGES

The major features of VHLLs are non-procedurality, high level operations and abstract data structures. A non-procedural description specifies a task in terms of its behavior independently of any specific way of accomplishing the task.

2.2.1 GENERAL PURPOSE VHLL

SETL[KESC 75] emphasizes non-procedural task specifications in terms of mathematical sets; APL[IVER 62] has many convenient high level operations on arrays. There are also special-purpose VHLLs being developed in the areas of simulation (SIMULA[DAMN 70], GPSS[BOKP 76]), and business data processing (SSL[NUNA 71], BDL[HHKW 77]). The non-procedurality of VHLLs presents problems of implementation and optimization which are more difficult than in High-Level Languages(HLL). This is because the choice of feasible execution algorithms must be made automatically. In addition, the abstract data structures requires the choice of suitable data representation also to be made automatically.

The programming language SETL trys to ease the programming problem by using powerful operations on very general data structures such that the issues of problem

formulation can be separated from those of program efficiency. Sets and tuples as well as other primitive data entities can be manipulated in the SETL language. Existential quantifier and universal quantifier can be used to construct a boolean expression similar to predicate calculus. In addition, universal quantifier can be used to form a loop over the elements of set entities such that the knowledge of data representation of sets is not necessary in describing the algorithm.

Program optimization is particularly important in VHLLs and there are many techniques that can be applied to improve efficiency. For example, the data structures of sets and tuples are not specified by the user in a program written in SETL. It may be a bit vector or a linked list or something else. The simplest translation of such a language will yield very inefficient programs. For this reason the need to optimize a program written in a VHLL is especially important. Also, the information that an optimizer needs is much more accessible in the abstract, problem-oriented specification of a VHLL than in the detailed code sequences of a language of lower level.

A non-procedural language LUCID[ASWA 77] has been designed as a formal system in which programs can be written and their proofs carried out. The statements of a LUCID program can be interpreted as true mathematical assertions about the results of the program. For example, an

assignment statement in LUCID can be considered as a statement of identity, or equation. A variable in LUCID has a history which is an infinite sequence of data objects. Special functions FIRST and NEXT can be used to reference the first element and the sequence starting from the second element of the history of a variable respectively.

In general, a LUCID program defines the histories of a set of variables by relating their histories with a set of equations. The use of FIRST and LAST functions allow basically the specification of one level loops. In order to allow nested loops, a function LATEST is introduced. It clutters up the program; consequently, BEGIN-END blocks to nest iterations are introduced into the language.

Although MODEL is not a language intended for automatic program verification, the spirit of the language is similar to that of LUCID in that the computations are specified with non-procedural mathematical assertions. In 1973, Ramirez used a data definition language[RAMI 73] as a tool to generate data conversion program automatically. Although the aim of his research was to save programming work in a special application, the concept of using data and computation descriptions to specify data processing tasks generally was introduced. Rin extended the work of Ramirez and developed an initial version of a non-procedural programming language called MODEL, limited to use in business transactions processing[RIN 76]. For each

transaction processing program, the programmer had to describe only the structure of input and output files and assertions describing relations between input and output data. The language processor analyzed the MODEL statements and generated a corresponding PL/I program. The programs generated by MODEL processor include: (1) proper input and output statements to get data in and out of the main memory and optionally some packing and unpacking statements if data is stored in variable format on external storage, (2) a list of assignment statements enclosed by very simple iteration control statements. The language processor analyzed the precedence relation between statements in a specification. For this purpose it used a directed graph. An executable program was generated from the graph.

Shastry considered MODEL as a general purpose language[SHAS 78]. He analyzed the subscript expressions occurring in array element references, where the subscript expressions could be first order polynomials. By the technique of splitting nodes in the graph, he transformed a cyclic graph into an acyclic one if the specification was sequenceable. He also conducted extensive analysis of consistency and completeness of the program specification to detect errors before the program was generated. Inconsistency could be due to invalid subscript range specification or due to inconsistent use of subscript names. Incompleteness could be due to the omission of the data

description statements for some data names or the omission of an assertion that defined a field of an output file. Any cycle in the array graph which corresponded to a set of simultaneous equations was considered not sequenceable.

The capability of automatic applying of numeric methods to solve a system of equations was incorporated into the MODEL processor[GREB 81]. It has proved useful in applications of econometric forecasting and modelling. Recent development of the MODEL system further extended the capability of the system. Modularity and execution of subspecifications in parallel or in distributed computation are currently under development. The proposition of extending the MODEL system for distributed computation is discussed in [PNPR 81]. The use of data flow computer to perform the computation in MODEL system is being explored by [GOKH 81].

One objective of use of VHLLs is to decrease the involvement of computer users in the complexity of computer characteristics. Although the introduction of HLLs has relieved programmers of the painstaking struggle with particular computer architectures, HLLs are still very far from the language that problems are discussed and solution methods are presented. Software development is still a laborious and difficult task to undertake. One of the approaches to ease the work of software development is through the use of VHLLs. VHLLs usually offer use of

abstract data structures, high level operations and non-procedurality. In this way the user can concentrate naturally on the problem statement without considering implementation related decisions that become entangled with the problem logic. In some cases the level of the languages is sufficiently high, requiring only a high level specification of the computations, which can be prepared by non-programmers.

It has been suggested that most of the conventional programming effort goes into selection of proper data representations and data manipulation algorithms to perform the computations efficiently[SCH 75]. Sometimes the consideration of program efficiency may cause the sacrifice of program readability and comprehension. In turn, it affects the ease of program testing and maintenance. The use of VHLLs offers many benefits such as less coding work, less required proficiency in programming and in algorithm analysis, and ease in understanding and updating the program. All these benefits are conditioned on whether the language processor can produce satisfactorily efficient programs.

Users of MODEL need not be concerned with physical representations of the data. MODEL processor allocates memory for each data structure in the specification. When all the elements along some dimension of an array can share the same program variable, we say that dimension of the

array is <u>virtual</u>. Otherwise, the dimension of that array is <u>physical</u>. Virtual array dimensions save memory space. In addition, users do not have to specify program controls such as loop control or I/O control.

Recently Rajeev Sangal [SANG 80] has investigated the possibility of introducing modularity in non-procedural languages such as NOPAL, a non-procedural language for automatic testing, and MODEL. The use of abstract data types is suggested as an approach to modularity. The abstract data types are specified in modules. A module consists of a header, data declarations for the representation of the abstract data type, and a set of module functions which are the allowed operations on the abstract data type. The functions are also defined within the framework of non-procedural languages.

2.2.2 PROBLEM ORIENTED VHLLS

Many problem statement languages have been developed to automate the system design of very large information systems. They allow the statement of requirements for an information system without stating the procedures that will be used to implement the system. The computer programs can be used to analyze the problem requirements and report the logical inconsistency and incompleteness to the system designer. For example, Accurately Defined Systems(ADS), a

product of the National Cash Register Company[LYNC 69], consists of a set of forms and procedures for a systematical approach to the system definition. An ADS requirements statement includes the descriptions of (1) inputs to the information system, (2) historical data stored by the information system, (3) outputs produced by the information system, and (4) actions required to produce these outputs and the conditions under which each action is performed. The ADS Analyzer can perform a number of checks, ranging from simple syntax checking to more complex logical consistency and completeness checking. It also produces a number of summary reports such as a dictionary of all data element occurrences, indices to all data elements and processes, data dependency matrices and precedence relationships among data elements and processes, and graphical displays of the ADS forms. The use of ADS can save the system designer considerable time during the specification of logical system design because the ADS Analyzer can provide them feedback before the physical design or coding starts.

SODA Statement Language(SSL) was developed by Nunamaker[NUKO 76]. It is designed for the total design process from non-procedural problem statement through software design and hardware selection to final implementation and performance evaluation. An SSL problem statement is composed of a collection of Problem Statement

Units(PSU). A PSU consists of three components: data description, processing requirements, and operational requirements such as information on volumes, frequency of output, and timing of input and output. The problem statement analyzer finds the precedence relationships between the data and processes, then uses the matrix algebra and graph theory to check the consistency and completeness of the problem statement. Another program called SODA/ALT determines the number of CPU and the size of core memory in the hardware system under the constraints of operational requirements. It then selects a program module and file design from feasible alternatives with the concern of reducing the total transport volume by grouping operations into modules and data sets into files.

Business Definition Language(BDL) is a very high-level programming language used in the domain of business data processing. The concepts in BDL were derived from mimicking a model of business organization. For example, the documents in BDL, which serve as input and output to a program as well as internal representation of information, correspond to the business form; steps in a program correspond to the organizational units of the system being described. In a Form Definition Component, the user defines the format and structure of the forms used in the program. The Document Flow Component is used to describe the interconnections of the steps in the same way as that used

to describe the business organization. The computations on the documents are described in the Document Transformation Component. The documents are routed among the various units of the organization or stored in files and computations on the elements of forms can be done in the basic steps.

The Requirements Language Processor(RLP)[DAVI 79] developed at GTE Laboratories aimed to automate the requirements phase of the software development. It is a table-driven compiler which allows the requirements to be written in a language that is designed specifically for the application area of the product. The RLP will accept the requirements of the system as input, produce formatted documents, report any incompleteness, inconsistency, ambiguity and redundancy in the requirements, and finally create a machine readable model of the specified system which is in the form of a finite-state machine. The FSM system model generated by the RLP can be used to help automate latter phases of software development [DAVI 80]. For example, the customer can apply a Feature Simulator over the system model to verify the system's behavior before design or implementation is initiated. Furthermore, a Test Plan Generator and an Automatic Test Executor can be used to automate the certification testing of the system based on the system model [BAFI 79].

2.2.3 VHLL OPTIMIZATION

In a very high-level language such as SETL, programs are written in terms of general data structures and their related operations. The compiler has to select the internal data representation and decide on the efficient algorithm to implement those high level operations. The optimization on this level emphasizes algorithm optimization which may have very significant effect on program execution and therefore is essential to the practical use of the language.

The design of very high-level languages emphasizes ease of use rather than efficient implementation. They usually allow use of high level operations on abstract data structures. However, the compilers have to translate high level operations into corresponding lower level operations and select data representations for abstract data structures. There may be many alternative algorithms that can be used to implement a high level operation. As is known, no amount of code optimization can compensate for a bad algorithm. The difference in performance between a clever and a naive program implementation can be quite significant. Therefore, optimization techniques applied to languages are essential if large programs written in these languages are to be run routinely.

In the language SETL, the objects being manipulated include finite sets, ordered n-tuples, and sets of ordered

n-tuples usable as mappings. It is the responsibility of the compiler to choose both the data structures which will represent the abstract objects in a program and the corresponding code sequences which will realize the abstract operations to be performed on these objects. For practical reasons, the choice is typically limited to the most representative data structures and the criteria which influence the choice of data structure are collected through an empirical study of manual translation. The optimizer performs global program analysis to check whether the criteria are satisfied.

Since the objects manipulated in SETL programs tend to be very complex data structures, it is desirable to pass a pointer rather than physically copy the data when an object is assigned to or made part of another variable. The SETL language takes value semantics for the assignment operation, i.e. the effect of assignment is to physically transfer some value from a source to a target variable instead of renaming the object being assigned as in CLU. This may cause problems in modification to the existing objects. The cases where a minor change to an existing object can be safely accomplished by modifying that object is discussed in [SCH 75]. Another major issue in optimizing a SETL program is to properly select the data structure. The decision may be based on the relationships of inclusion and memberships between objects in the program. The technique to discover

these relationships is described in [SCHW 75].

In a business-oriented automatic programming system such as PROTOSYSTEM-I and SODA the optimization concentrates on the reduction of number of I/O accesses. The method to reduce the number of accesses is through merging of data sets and computations. By aggregating the data sets which have the same key field into one physical file, many related data items can be accessed from a single data file when they are needed for processing, rather than having to access them from several different files. There are two ways to aggregate computations such that the number of accesses can be reduced. When several computations require the same input data sets, it is desirable to group all of them into one computation. The benefit is that a record to be accessed need be read once for all the computations, rather than once for each computation. The aggregation of two computations may be advantageous when the output of one is fed as the input to the other. In this case, the need for the latter computation to read output records of the former is eliminated. If the output of the former computation is not further used by any other computations, the writing out of the data set can be eliminated, too.

In the MODEL system, programs are optimized by selecting efficient loop control and memory allocation schemes based on a non-procedural specification. A part of the program design module has knowledge about what

alternative loop structures are feasible to implement the required computation and another part of the module will evaluate the quality of each alternative in terms of the overhead of loop control statements and the amount of memory space for program data. A phenomenal program improvement can be achieved by maximizing the loop scopes in the program. The consideration of merging two loops is not limited to the case that they iterate same number of times. When the instances of one loop correspond to a subset of those of another loop, we may still merge the two loops into one. This feature of allowing loops with different number of iterations to be merged makes the efficient implementation of list like data manipulations possible. Although the optimization techniques that we have developed are used primarily for the MODEL system, with some preprocessing it is possible to apply them to other array-oriented VHLL such as APL. For APL, the necessary preprocessing is to rename the program variables when the same variable names are served for different uses such that an APL program will become a non-procedural program specification. After an APL program has been transformed into a program specification, it can be submitted to the MODEL system to generate an efficient program.

2.2.4 SOURCE-TO-SOURCE TRANSFORMATION

Some systems perform a source-to-source transformation on the program representation to improve or refine a program. The motivation for the program transformation systems is to encourage users to write programs which are easy to read, understand, and update, without having to consider program efficiency. These programs are transformed in a systematic way into a more intricate but effiecient form.

From the view point of ease of program maintenance, programmers should be encouraged to write programs that are easy to read and easy to change. It is advisable, therefore to adopt a discipline in the programming style. However, such a program may suffer a heavy penalty in program running time. In practice, it is often necessary to trade program comprehensibility for program efficiency. The technique of source-to-source transformation aims to overcome this dilemma by manipulating a program in its source representation into an efficient version.

Early attempts of source-to-source transformation made the program improvement visible to the user [SCAN 72]. Optimizing programs at the source level usually also requires that the optimization techniques are machine independent. Some of the program transformation system emphasize program optimization and others emphasize program

refinement.

Burstall and Darlington[BUDA 77] described a system which can convert program structure from recursion to iteration and transform data structures from abstract to concrete. The program to be transformed is presented as a set of recursion equations. Transformations rules such as definition, instantiation, unfolding, folding, and abstraction can be used to add new definitions of functions into the set. Heuristic strategies for applying the transformation rules are used to help avoid fruitless search. The process of producing new definitions for functions continues and hopefully the more efficient versions of the function definition will be generated by the system. The same program transformation technique can also be used to help abstract programming. The user is required to define a single representation function which maps the lower data type onto the higher, then programs written in terms of higher level primitives can be rewritten in terms of the lower level primitives by the system.

The Program Develpoment System (PDS) developed at Harvard University aimed to simplify the work of program maintenance [CHTH 79] [CHHT 81]. The system takes an abstract algorithm as input and applies a set of user-defined transformation rules to the abstract algorithms, then produces an efficient program which realizes the algorithm. Since the implementation decisions

which are program efficiency relevant can be incorporated in the user-defined transformation rules, programs can be designed and modified in their abstract forms. The same program efficiency considerations will be maintained by applying the program transformation again. A transformation rule consists of a syntactic pattern part, optionally augmented by a semantic predicate, and a replacement part. Since both the program to be transformed and the transformation rules are converted to a tree representation, the transformation process is basically subtree matching and replacement.

Two classes of program transformation techniques discussed by Kuck[KKPL 81] aim to transform FORTRAN programs into a form which exploits the computer architecture capable of parallel processing. A collection of techniques based on simple rewriting transformations remove unnecessary dependency relationships between program statements. When a program is to run on a machine with parallel processing capability, reducing the number of dependencies usually leads to a reduction in the program's running time. Sharing the same variable for different values is adequate for sequential programs. However, it imposes unnecessary sequentiality constraints on parallel programs. The renaming transformation which assigns different names to different uses of the same variable and the expansion transformation which changes a variable used inside a loop

into a higher dimensional array remove the sequentiality constraints caused by sharing the memory space. Another class of transformation aims to reconfigure the loop structures in a program such that the scope of recurrence loop is reduced and the possibility of doing vector operations is increased, which in turn speeds up the execution. A technique called loop distribution breaks loops into smaller ones as long as possible. On the other hand, in a virtual memory environment merging two loops which reference the same set of vectors is helpful to reduce unnecessary page swap.

In order to facilitate further the use of the MODEL language in the areas of mathematical computation and data processing, operations on higher level data structures and matrix operations are proposed as an extension to the system. The technique of source-to-source transformation has been studied for implementing those features. A statement containing high level operations is replaced automatically by a set of statements containing only lower level operations. This extension essentially increases the level of abstraction in specifying computations and potentially reduces the number of mistakes made by the user.

2.3 AUTOMATIC PROGRAM SYNTHESIZING SYSTEMS

Automatic programming systems usually synthesize programs from problem specifications in particular application domains. They can be divided into the knowledge-based approach and the formal-model-based approach. Knowledge-based automatic programming systems such as PSI[GREE 77] and OWL[SZHM 77] contain a great deal of information about some application domain. They accept very high-level problem descriptions, check for consistency and completeness, and use knowledge about the application domain to translate the problem description into a procedural program which satisfies the problem requirement. Formal-model-based automatic programming systems such as PROW[WALE 69] derive program from logic theorem proofs. They accept the problem specification and the primitive operations in the form of logic formulas. Then the theorem proving techniques are used to synthesize the required programs.

PSI is a knowledge-based automatic programming system developed at Stanford University. It consists of a set of closely interacting modules or experts. A discourse expert is responsible for conducting a dialogue with the user in natural language. A domain expert interprets terms with domain-specific meanings and provides help to both the user and the model-building expert regarding possible algorithms and information structures to be used. A trace

expert[PHIL 77] allows the user to specify a program with the trace of the program execution. The model-building expert[MCCU 77] contains high-level general programming knowledge and rules for assemblying fragments of program description coming from the domain expert into a complete program model. After the program model is built up, it is passed to the coding expert[BAKA 76] which produces an efficient target language program with the help of the efficiency expert.

The synthesis phase of the PSI system constructs programs from high level program models with a coding expert and an efficiency expert. The coding expert uses rule-based programming knowledge to produce alternative algorithm and data structure choices. The program optimization is performed by the efficiency expert which estimates space-time costs for every partially developed program passed from the coding expert[BAKA 76]. The estimation is performed with an exact mathematical analysis on the number of times that each statement is executed. For statements within loops, the efficiency expert computes the average number of executions by summing the probability of execution over all possible loop instances. The branch probability of a conditional test and the execution probability of a loop instance which are essential to the estimation of execution frequency are either assumed by the efficiency expert or from user's comment. For every statement in the partially

developed program, the efficiency expert computes its execution frequency, space usage, and single execution time. Then the space-time product is used as the cost function. The alternative with the smallest cost will be picked as the best choice.

The OWL system is the top-part of a automatic program generation project at MIT. It aims to be a knowledge-based man-machine interface which can accept the problem description in natural language and produce a data processing specification. Its application domain is in the area of Management Information Systems. The bottom part of the project, PROTOSYSTEM-I, obtains a problem statement written in SSL from the top part. It analyzes the specification, performs the system design, and generates PL/I code and JCL for the required system.

The formal-model-based automatic programming system started with the idea of deriving programs automatically with a mechanical theorem prover. [GREE 69], [MANN 71], [LEWA 74] In order to construct a program, the user first formulates the relation between the input and the output variables of the program. Then the system proves a theorem induced by this relation and extracts the program from the proof directly. Since the program is derived form its logical specification, it does not require debugging or verification. For example, the PROW system by Waldinger and Lee accepts the specification of a program written in the

language of predicate calculus, decides the algorithm for the program, and then produces a LISP program which is an implementation of the algorithm. The instructions of LISP are axiomatized and stored as axioms in PROW. The input and output relationship of the program is expressed as a well-formed formula in the first order predicate calculus. A logic theorem is constructed from the program specification and a theorem prover is invoked to generate a proof of the theorem. The desired program is then extracted from the proof of the theorem.

CHAPTER 3

SYNTAX AND SEMANTICS OF THE MODEL LANGUAGE

3.1 STRUCTURE OF A PROGRAM SPECIFICATION

A program specification written in the MODEL language consists of three major parts: program header, data description, and assertions. The program header specifies the name of the program and the external files which store the input or output data of the program. The data description statements are used to specify the data structure of the input or output files and the structure of the intermediate results. The assertions are used to define the values of the intermediate or output variables specified in the data description statements. Although the user is encouraged to group statements together and order the parts in the sequence mentioned above, the statements in a program specification can be put in any order, i.e. the order of the statements is irrelevant to the meaning of the specification. That is one reason why we call MODEL a

non-procedural programming language. In this section we discuss the statements in the program header. We will discuss in section 3.2 the data description statements, and in section 3.3 the syntax and the semantics of the assertions. We will discuss in section 3.4 the use of control variables.

Only the basic MODEL language is described here. Short-hand and high level dialects are not described as they are always translated automatically into the basic language. The syntax rules of the MODEL statements will be defined with extended BNF notation. Identifiers enclosed by the angle brackets ('<' and '>') are non-terminal symbols. The metasymbols used include:

1. ::=, it is read as 'is-defined-by'.

2. [...], a pair of square brackets is used to enclose a string which is optional.

3. |, a vertical bar is used to separate alternatives.

4. {...}*, a pair of braces followed by an asterisk is used to enclose a string which can repeat any times (including zero).

The program header consists of three types of statements, namely the <u>module</u> statement, the <u>source</u> <u>file</u> statement, and the <u>target</u> <u>file</u> statement.

Module Statement

The syntax rule for the module statement is as follows.

<module-statement>::=

    MODULE : <identifier> ;

The user-chosen identifier is used as the name of the program being specified.

Source File Statement

The syntax rule for the source file statement is as follows.

<source-file-statement>::=

    SOURCE [FILES | FILE] : <identifier> { , <identifier> }* ;

The source file statement consists of a list names of files which serve as the input files of the program. The source files are assumed stored in external storage devices.

Target File Statement

The syntax rule for the target file statement is as follows.

<target-file-statement>::=

    TARGET [FILES | FILE] : <identifier> { , <identifier> }* ;

The target file statement lists the names of files which serve as the output files of the program. The output files are assumed to be on external storage and they serve

to retain the computation result for future use.

3.2 DATA DESCRIPTION STATEMENTS

In a non-procedural programming language every variable can only have a single value. Therefore, different variable names should be declared for different data involved in the computation. The data structures in external files, or the schemata of files, can be described in MODEL with data description satatements. Logically related variables may also be grouped together as in PL/I. The user must also declare the data types of the components of a variable in data description statements. The MODEL language has been designed to relieve the user of concern for I/O control. In general, I/O can be a complicated part of a programming language. A few simple mechanisms have been included in the data description statements to ease the I/O programming task. Examples include the ability to describe file organization and to indicate a key field for direct accessing a record. In section 3.2.1 we will discuss the way to specify the data type of a variable; in section 3.2.2, the way to describe data aggregates; and in section 3.2.3, the mechanisms used for I/O related programming.

3.2.1 DATA TYPES

The smallest unit of data in a program is a field.   A
field may contain a datum of some type supported by the
MODEL language. The available data types includes picture,
character, bit string, and numbers.   It is the user's
responsibility to select a data type for each field.

Field Declaration Statement

The syntax rule for a field declaration statement is as
follows.

```
<field-declaration-statement> ::=

    <identifier> [ IS ] <field> <data-type> ;

<field> ::=  FLD | FIELD

<data-type> ::=  <type> <leng-spec>

<leng-spec> ::=  ( <min-length> [ : <max-length> ] )

<min-length> ::= <integer>

<type>::= <pic-desc> | <string-spec> | <num-spec>

<pic-desc> ::=  <pic-type> ' <string> '

<pic-type> ::= PIC | PICTURE

<string-spec> ::= CHAR | CHARACTER | BIT | NUM | NUMERIC

<num-spec> ::= <num-type> [ <fixflt> ]

<num-type> ::= BIN | BINARY | DEC | DECIMAL

<fixflt> ::= FIX | FIXED | FL | FLOAT | FLT

<max-length> ::= <integer>
```

A character string may be of fixed length or variable length. For a fixed length character string the length in byte units should be specified in the type declaration. A variable length character string is specified through declaring the range of the possible length of the string. When a field X of variable length string occurs in an input file, its length should be specified by an associated control variable called LEN.X.

Example:

```
A IS FIELD CHAR(6) ;
B IS FIELD CHAR(0:10);
```

The field A is a string of six characters and the field B is a variable length character string with maximum length ten. The actual length of the field B should be specified by a control variable called LEN.B in some assertion.

The available operations for manipulating character strings include lexicographic comparison, concatenation, and extracting substring. The discussion for the character string is also applicable to the bit string data type.

The data types for numeric data include picture, floating point decimal, floating point binary, fixed point decimal, and fixed point binary. The operations applicable to numeric data are arithmetic operations, comparison, and conditional definition. It should be noted that the picture and character typed variables have a printable representation. Therefore, it is suitable for data

contained in reports. Other numeric data types are generally used for the data stored in the computer system. The PL/I target language incorporate extensive type conversion and therefore the user is generally relieved of this concern.

3.2.2 DATA STRUCTURES

Usually there are two ways to group logically related data together to form data structure. An _array_ contains homogeneous data elements and a _structure_ contains heterogeneous data elements. In MODEL a generalized data aggregate can be used to specify arrays and structures. The data aggregate is called a _group_ or a _record_ in MODEL language.

Group Declaration Statement

The syntax rule for the group declaration statement is as follows.

```
<group-declaration-statement> ::=
    <identifier> [ IS ] <group> ( <member-list> ) ;
<group> ::= GRP | GROUP
<member-list> ::= <member> { , <member> }*
<member> ::= <identifier> [ ( <occspec> ) ]
<occspec> ::= * | <minocc> [ : <maxocc> ]
<minocc> ::= <integer>
```

`<maxocc> ::= <integer>`

In the group declaration statement an identifier is declared as a data group which contains a list of members. Each member may optionally repeat some number of times. If a member repeats, it is considered as an array of one dimension more than the group containing it. There are three ways to specify the number of repetitions over a dimension of an array. If the number of repetitions is a constant, then the constant can be specified along with the array name. When the number of repetitions is not fixed but the user knows the maximum of it, he can specify a range for the number of repetitions in the group statement. If the user does not know the maximum, i.e. where the maximum is an unknown large value, he can denote the range by an asterisk. When the number of repetitions is not a constant, it can be defined through some control variables with keyword prefix such as SIZE or END (refer to section 3.4) or definition may be omitted if it can be detected based on an end-of-file indication.

The members of a data group can be fields, or some other data groups. A data group may be declared as an array of arrays. In order to reference a unit datum of it, the user has to supply as many subscripts as the number of array dimensions. Thus the member field becomes a multi-dimensional array.

Example:

```
A IS GROUP (B, C(10)) ;

B IS FIELD CHAR(6) ;

C IS GROUP (D(5), E(1:50), F(*)) ;
```

where identifier A is declared as a data group containing two members B and C. Let us assume that A is a zero dimensional variable. Since C repeats, it is a one dimensional array. Identifier C contains three members, D, E, and F. The member D repeats five times, and the member E may repeat a number of times from one to fifty. The member F has a unknown number of repetitions, so an asterisk is specified as its number of repetitions. All the members of data group C are two dimensional arrays.

3.2.3 I/O RELATED DATA AGGREGATES

In a MODEL specification, the user describes the structures of the data files with data description statements. The MODEL processor generates I/O statements automatically for the source and target files of the program based on the information in data description statements.

The record declaration statement is syntactically similar to the group declaration statement. The only difference is that the keyword GROUP is changed to RECORD. A record corresponds to a unit of data which can be

physically transferred between external file and main memory.

The file is the highest-level data structure which could be declared in a MODEL specification. It is not allowed to have a structure above the file. A file structure may consist of substructures declared with group, record, or field statements. A well structured file declaration will have the file entity on the top level. Its immediate descendants (i.e. members) can be declared either as groups or records. The groups may contains groups, records, or fields. Finally on the lowest level in the file structure the data should be declared as fields.

File Declaration Statement

The syntax rule for the file declaration statement is as follows.

<file-declaration-statement> ::=

    <identifer> [IS] FILE [NAME] <file-desc>

     (<member-list>) ;

<file-desc> ::=

    [KEY [NAME] [IS] <identifer>]

    [ORG [IS] <org-type>]

<org-type> ::= SAM | ISAM

A file may have the KEY attribute specified. In that case, the records in the file are accessed by a part of the record contents. If a file is keyed, there can only be one

record type in the fi'. structure and one of the field in
the record should be c..-ared as the key for accessing the
record. Two types of file organization are supported by the
MODEL language, namely the sequential files and the index
sequential files. A record in an index sequential file can
be accessed faster than in a sequential file if direct
accessing is necessary.

Example:

```
MODULE: MINSALE;
SOURCE: TRAN, INVEN;
TARGET: SLIP, INVEN;

TRAN IS FILE (SALEREC(*));
    SALEREC IS RECORD (CUST$,STOCK$,QUANTITY);
        CUST$ IS FIELD(CHAR(5));
        STOCK$ IS FIELD(CHAR(8));
        QUANTITY IS FIELD(CHAR(3));
INVEN IS FILE (INVREC)
        KEY STOCK$
        ORG ISAM;
    INVREC IS RECORD(STOCK$,SALPRICE,QOH);
        STOCK$ IS FIELD(CHAR(8));
        SALPRICE IS FIELD(NUMERIC(5));
        QOH IS FIELD(NUMERIC(5));

SLIP IS FILE (SLIPREC(*));
    SLIPREC IS RECORD (CUST$,STOCK$,QUANT,PRICE,CHARGE);
        CUST$ IS FLD (CHAR(12));
        STOCK$ IS FIELD(CHAR(16));
        QUANT IS FIELD (PIC'(11)Z9');
        PRICE IS FIELD (PIC'(11)Z9');
        CHARGE IS FIELD (PIC'(11)Z9');
```

3.3 ASSERTIONS

Data description statements define the data structures
of the variables involved in a computation. However, the

values of the variables are defined either automatically by input files or manually by assertions. Basically an assertion is an equation. On the left hand side of the equal sign there should be either a simple variable or a subscripted array name which references an array element. On the right hand side there can be any arithmetic or logical expression whose value is used to define the variable on the left hand side. The current restriction is that the assertion can only be used to define the value of a field. Operations on the higher level data structures are proposed to be translated into basic operations [PNPR 80].

3.3.1 SIMPLE AND CONDITIONAL ASSERTIONS

There are two kinds of assertions which can be used to define the value of a variable, namely <u>simple assertion</u> and <u>conditional assertion</u>. The assertions have the same syntax as an assignment statement and a conditional statement in the PL/I language, respectively. All the arithmetic and logical operations can be used in composition of expressions. In addition, the conditional expression of ALGOL language can be used in composing the expression.

<u>Simple Assertion</u>

The syntax rule for the assertion is as follows.
<assertion> ::= <simple-assertion> | <conditional-assertion>

```
<simple-assertion> ::= <variable> = <expression> ;
<variable> ::= <simple-variable> | <subscripted-variable>
```

The variable name on the left hand side of an assertion is called the **target variable** of the assertion as its value is defined by the assertion. All the variables on the right hand side are called the **source variables** of the assertion since their values are used to calculate the value of the target variable. In the examples shown below, a conditional expression is used to define the value of variable M.

Example:

    1) A = B + 5 ;

    2) X(I,J) = 4 * I + J ;

    3) M = IF OK THEN 5 ELSE 0 ;

Conditional Assertion

The syntax of the conditional assertion is similar to that of an IF statement in PL/I.

```
<conditional-assertion> ::=
    IF <boolean-expression> THEN <assertion>
                    [ ELSE <assertion> ]
```

The conditional assertion has two branches, one after the keyword THEN and the other after the keyword ELSE. These two branches are selectively executed according to the truth value of a boolean expression. Since the purpose of an assertion is to define the value of a variable, there can only be one target variable in an assertion. In any case

the two branches should define the same target variable. Therefore, the target variable in any branch of a conditional assertion should always be the same. It should be noted that the ELSE branch of a conditional assertion is optional. If it is omitted, the target variable may be undefined in some cases.

Example:

```
1) IF I < 5 THEN A(I) = B(I) ;
         ELSE A(I) = B(I) + 2 ;
2) IF END.X(J) THEN B = X(J) ;
```

3.3.2 SUBSCRIPT EXPRESSIONS

The variables used in assertions are either simple variables or subscripted variables. A specific element of an N dimensional array can be referenced with the array name followed by N subscript expressions. In the following we will discuss how the subscript expressions are formed and how they are used in composing the assertions.

Subscript expressions are composed of ordinary variables, subscript variables, and constants with arithmetic operations. The subscript variable is a special kind of variable. It does not have structure and it does not hold one specific value. Instead, a subscript variable assumes integer values in a range from one up to some

positive integer. If the range for a subscript variable is fixed in the whole program specification, then the subscript variable is called a global subscript. On the other hand, if the range for a subscript variable is to be determined for each assertion, the subscript variable is called a local subscript. There are ten system predefined local subscripts named SUB1, SUB2, ..., up to SUB10. There are two types of global subscripts. One of them has the form of qualifying the name of a repeating data structure prefixed with the keyword FOR_EACH. The other is created by declaring an identifier as a global subscript with the subscript statement.

Subscript Declaration Statement

The syntax rule for the subscript declaration statement is as follows.

<subscript-declaration-statement> ::=

    <identifier> IS <subscript> [(<occspec>)] ;

<subscript> ::= SUBSCRIPT | SUB

The subscript expressions are classified into the following types according to their forms. In the following, let I denote a subscript variable, c and k denote non-negative integers, and X denote an indirect indexing vector(refer to section 4.2.2.2.) Subscript expressions may be classified as follows:

    1) I,

2) $I-1$,

3) $I-k$, where $k>1$,

4) none of the other types,

5) $X(I)$

6) $X(I-c)-k$, where $c+k=1$,

7) $X(I-c)-k$, where $c+k>1$.

The range of a global subscript variable in an assertion may be declared in a subscript declaration statement. If not declared, the range is derived from an array dimension in which the subscript variable has been used in a type 1, 2, or 3 subscript expression.

Example:

1) I IS SUBSCRIPT (10) ;

   $B(I) = A(I)$;

   A global subscript I is declared in the subscript declaration statement and the range of the value of I is from one to ten. In the assertion, the global subscript I will assume the integer values in the range declared in the subscript declaration statement.

2) FACT(SUB1) = IF SUB1=1 THEN 1

                          ELSE SUB1 * FACT(SUB1-1) ;

   The range of the local subscript SUB1 will be the same as that of the first dimension of array FACT because the subscript SUB1 occurred in the term FACT(SUB1) is in a form of type 1 subscript expression.

The use of subscript variables allows us to define all the elements of an array in one assertion. In the second example above, the whole vector FACT is defined by the same assertion.

For multi-dimensional arrays, subscripting array variables may become tedious. We have adopted the following convention to allow users to omit subscripts in array references. When all the array references in an assertion have the same leftmost subscript expression, which is a type 1 subscript and when the subscript is not otherwise referred to in the assertion, then the subscript can be omitted from the assertion systematically. For example, the following three assertions are equivalent.

    a1: A(I,J,K) = 2 * B(I,J,K) + C(I,J) ;
    a2: A(J,K) = 2 * B(J,K) + C(J) ;
    a3: A(K) = 2 * B(K) + C ;

3.4 CONTROL VARIABLES

Sometimes it is necessary to refer to attributes of the data, such as the number of repetitions, the length, or the key for accessing a record in an index sequential file. In order to allow reference to such attributes, a number of control variables are included in the MODEL language. Since the control variables are always related to some variable, they have a form of a qualified variable, with the name of

the variable as the suffix and one of several reserved
keywords as the prefix. In the following we will assume
that X is a variable name declared in some data description
statement. The control variables which can be formed from X
are discussed below.

SIZE.X

If X is a repeating member of some data structure, the
user can specify the range by defining the value of a
control variable called SIZE.X. It should be noted that X
may be a multi-dimensional array. SIZE.X defines only the
range of its rightmost dimension. The ranges of the other
dimensions have to be defined separately.

SIZE.X is a variable of integer type. Its value is
used to specify the number of repetitions of the rightmost
dimension of array X. If $X(I1,I2,...,In)$ is an n
dimensional array where I1 occurs on the most significant
dimension and In on the least significant dimension, then
the control variable $SIZE.X(I1,I2,...,Ik)$ should be a k
dimensional array with $0<=k<n$. The first dimension of
SIZE.X has the same range as the first dimension of array X,
the second dimension has the same range as the second
dimension of array X, and so on. The value of SIZE.X cannot
be a function of any subscript Ii with $k<i<=n$. For every
n-1 tuple $(I1,I2,...,In-1)$ which corresponds to a possible
combination of the leftmost n-1 subscripts for array X, the

number of elements of array X with this tuple as their leftmost n-1 subscripts is specified by the array element SIZE.X(I1,I2,...,Ik).

Example:

```
A IS GROUP (B(3)) ;
B IS GROUP (C(*)) ;
C IS FIELD ;
SIZE.C(1) = 4 ;
SIZE.C(2) = 2 ;
SIZE.C(3) = 3 ;

SIZE.C         C
-----          ------------------------------------------
| 4 |          | C(1,1) | C(1,2) | C(1,3) | C(1,4) |
-----          ------------------------------------------
| 2 |          | C(2,1) | C(2,2) |
-----          ----------------------------
| 3 |          | C(3,1) | C(3,2) | C(3,3) |
-----          --------------------------------
```

In the example above, array C is two dimensional. There are three instances of B in data group A and each instance of B contains a number of elements of array C. Correspondingly the range of the first dimension of array C is a constant three and the range of the second dimension which may depend on the subscript value of the first dimension is specified in vector SIZE.C. SIZE.C(1) equals to four implies that there are four elements of array C in the first instance of B, the value of SIZE.C(2) specifies the number of elements of array C in the second instance of B, and so on.

END.X

If X is a repeating member of a data structure, END.X can be used to specify the range of the rightmost dimension of array X as alternative to the use of SIZE.X.

END.X is a boolean array. If $X(I1,I2,...,In)$ is an n dimensional array, then the associated control array $END.X(I1,I2,...,In)$ is an n dimensional array, too. The range of array dimensions of END.X are the same as the corresponding array dimensions of X. The value of END.X determines the range of the rightmost dimension of array X in the following way. For every n-1 tuple $(I1,I2,...,In-1)$ which is a possible combination of the leftmost n-1 subscripts of array X, there exists a sequence of elements in END.X array with the same left n-1 subscript values, i.e. $\{END.X(I1,...,In-1,In)|\ 1<=In\}$. If $END.X(I1,...,In-1,m)$ is a boolean true and all the elements of $\{END.X(I1,...,In-1,In)|\ 1<=In<m\}$ are false, then there are exactly m elements in array X with $(I1,...,In-1)$ as their leftmost n-1 subscripts. The values in END.X may depend on the values in array X, i.e. the number of repetition may depend on the data in X.

Example:

For the same array C mentioned above, we may use a two dimensional control array END.C to specify the range of the second dimension of array C as follows.

    A IS GROUP (B(3)) ;

```
B IS GROUP (C(*)) ;
C IS FIELD;
END.C(SUB1,SUB2) = IF SUB1=1 THEN (SUB2=4)
                   ELSE IF SUB1=2 THEN (SUB2=2)
                   ELSE IF SUB1=3 THEN (SUB2=3) ;
```

```
        C
          ------------------------------------
          | C(1,1) | C(1,2) | C(1,3) | C(1,4) |
          ------------------------------------
          | C(2,1) | C(2,2) |
          -------------------------
          | C(3,1) | C(3,2) | C(3,3) |
          ----------------------------

    END.C
          ------------------------------------
          |   F    |   F    |   F    |   T    |
          ------------------------------------
          |   F    |   T    |
          -------------------------
          |   F    |   F    |   T    |
          ----------------------------
```

In the first row of END.C the first boolean true comes in the fourth element, therefore, the fourth element is the last element in the first row of array C.   Similarly, the second element of the second row of END.C is true implies that there are only two elements in the second row of array C.

Example:

We will show how the END control variable can be used to specify a varying number of repetitions by finding the greatest common divisor of two positive integers M and N. Euclid's algorithm is used here.

```
MODULE: TEST ;
SOURCE: IN ;
TARGET: OUT ;
```

```
IN IS FILE (INR) ;
    INR IS REC(M,N) ;

OUT IS FILE (OUTR) ;
    OUTR IS REC(GCD) ;

WK IS GROUP (WKG(*)) ;
  WKG IS GROUP (WK1,WK2) ;
(M,N,GCD,WK1,WK2) IS FIELD NUM(4) ;

WK1(SUB1) = IF SUB1=1 THEN M
            ELSE IF WK1(SUB1-1)>WK2(SUB1-1) THEN
                      WK1(SUB1-1)-WK2(SUB1-1)
            ELSE WK2(SUB1-1) ;

WK2(SUB1) = IF SUB1=1 THEN N
            ELSE IF WK1(SUB1-1)>WK2(SUB1-1) THEN
                      WK2(SUB1-1)
            ELSE WK1(SUB1-1) ;

END.WKG(SUB1) = WK1(SUB1)=WK2(SUB1) ;

IF END.WKG(SUB1) THEN GCD = WK1(SUB1) ;
```

POINTER.X

If X is a record of a keyed input file F, the instances
of the record X can be selected and ordered according to the
value of a control variable POINTER.X. The control variable
POINTER.X has the same number of dimensions and the same
shape as the array X.   For every value in the control
variable POINTER.X, a record instance in the file F with
that key value will be presented in the corresponding
element of array X.   In order to use POINTER control
variable for selecting and ordering the records in a keyed
file, one of the field in records should be declared as a
key in the file declaration statement. The content of the
POINTER control variable is used as the key to access the

corresponding record from the keyed file.

A keyed file may either have sequential or index sequential organization. If the file is index sequential, the records stored in the file may be in any order. However, if the file is actually a sequential file, then the records have to be sorted in an ascending order according to the key field and the keys used to access the records should also be sorted in the same order. This is an implementation restriction. Without this restriction we can not read all the records we want from that file in one pass.

When a keyed file is declared as a source and a target file, the target file will be an updated version of the source file. Effectively only the records being selected may be modified. For the rest of the file they are kept intact in the target file. This mechanism makes the update of sequential or index sequential file much easier to specify. Since a key value may occur more than once in the POINTER array, the corresponding (one) record will be accessed, possibly updated, and written out several times. In order to make sure every update to the same record is effective, the updates have to be done sequentially. We can envisage that a new version of the keyed file is created after one record is updated and every update is done on the most recent version of the file.

Example:

In the following MODEL specification a source file INVEN is declared as a keyed file. STOCK$ in the record INVREC is the key field of INVEN file. Since the control variable POINTER.INVREC is equal to the field STK in file TRAN, the INVREC records will be ordered according to the values in the STK field.

```
MODULE: MINSALE ;
SOURCE: TRAN, INVEN ;
TRAN IS FILE (SALEREC(*)) ;
    SALEREC IS RECORD (CUST$,STK,QUANTITY) ;
        CUST$ IS FIELD(CHAR(5)) ;
        STK IS FIELD(CHAR(8)) ;
        QUANTITY IS FIELD(CHAR(3)) ;

INVEN IS FILE (INVREC(*))
  KEY STOCK$
  ORG ISAM ;
    INVREC IS RECORD(STOCK$,SALPRICE,QOH) ;
        STOCK$ IS FIELD(CHAR(8)) ;
        SALPRICE IS FIELD(NUMERIC(5)) ;
        QOH IS FIELD(NUMERIC(5)) ;

POINTER.INVREC = TRAN.STK ;
```

FOUND.X

If X is a record in a keyed file, then it is accessed through the value of a POINTER control variable. It may happen that the key value used to access the record does not match with any record. The accessing would fail. The user may test the value in a control variable called FOUND.X to find out whether a record with some specific key exists or not. This informaton may be used to decide whether a new record should be added into the file or an old record should be updated. The control variable FOUND.X has the same shape

as array X and POINTER.X. Its data type is boolean.

LEN.X

If X is a field in some record and its data type is variable length character string, then the actual length of X is specified by the control variable LEN.X which is used to disassemble the input or output records. Corresponding to every element of array X, there is an element in LEN.X. The values in the array LEN.X are integers. We can use any integer type expression to define LEN.X. The only restriction is that the content of LEN.X should not depend upon any data physically positioned in a record after the data field X.

NEXT.X

If X is a field in an input sequential file, the control variable NEXT.X can be used to denote the same field in the next physical record of the file. Although the next record usually means the record with a subscript value one larger than the current record, it may not be true when the current record is the last record in some group. The problem is caused by the fact that the user is dealing with structured data but the real data in the external file is in a linear form. Sometimes the information used to transform a sequence of records into a structured form can only be conveniently expressed in the way that the records are physically contiguous. For example, we may want to compare

the value of a key field in two adjacent records to determine whether a record is the last record in a group or not. The fact that the current record and the next record may or may not be in the same group causes trouble in referencing the next record.

Example:

Suppose the records in a transaction file contain a customer number and some relevant information and the records are sorted according to the value of the customer number field. We may use the following specification to describe the data structure.

```
TRANSACTION IS FILE (CUSTOMER(*)) ;
    CUSTOMER IS GROUP (TRANS_REC(*)) ;
        TRANS_REC IS RECORD (CUSTOM_NO,INFORMATION) ;
            CUSTOMER_NO IS FIELD (PIC'99999999') ;
I IS SUBSCRIPT ;
J IS SUBSCRIPT ;
END.TRANS_REC(I,J) =
        CUSTOMER_NO(I,J)^=NEXT.CUSTOMER_NO(I,J) ;
```

The term NEXT.CUSTOMER_NO(I,J) in the last assertion can not be replaced by CUSTOMER_NO(I,J+1) because there may not be a record with this pair of subscript values. The restriction in using the control variable NEXT.X is that the position of X field in a record should be fixed, i.e. the fields to the left of the field X can not be variable length strings or repeating with a variable number of times. Otherwise, the field X in the next record may not be located

correctly.

SUBSET.X

If X is a record in an output file, then the control variable SUBSET.X can be used to selectively omit some records from an output file. The SUBSET.X control variable is a boolean array of the same shape as the array X. When an element in the SUBSET.X has a value of boolean true, the corresponding record X will be put into the output file. On the other hand, if the element has a value of boolean false, the corresponding record will not be put into the output file. It should be noted that the use of SUBSET control variable does not affect any other computations. Only a subset of records X may be omitted from the output file.

CHAPTER 4

PRECEDENCE ANALYSIS

4.1 INTRODUCTION

A MODEL specification consists of many data description or assertion statements. In principle, the data description statements specify the structure of <u>data</u> <u>entities</u> such as file, group, record, and field. The assertions specify the relationships between the data entities. The data entities and the assertions are referred to here as <u>program</u> <u>entities</u>. On the other hand, in an executable program there are <u>program</u> <u>events</u> such as I/O activities, computations, or getting data ready. The events in a program generated by the MODEL system correspond to entities in the specification. For example, a file entity corresponds to an event of opening a file or closing a file; a record entity corresponds to reading a record or writing a record; and an assertion entity corresponds to computing a target variable. The sequence of the program events is not given by the user. Instead, it is determined by the MODEL processor under the

constraints of precedence relationships among the program events. In this chapter we discuss the analysis for recognizing the precedence relationships between program events and representing them in a directed graph.

Based on the specification we can find the unique symbolic names assigned by the user to data entities. Additionally the MODEL processor automatically assigns a unique name to every assertion. Similar to other compilers, the MODEL processor maintains a symbol table called dictionary which contains all the symbolic names of program entities and their attributes.

The dictionary is created by a procedure CRDICT which finds all the entities in the program specification and stores their names into the dictionary. Except for some special cases described below, there is a correspondence between each statement in the specification and an entity in the dictionary.

Attributes of a symbol such as the type (file, group, field, ..., etc), the number of dimensions, the structural relation of it to other symbols are stored in the dictionary during the process of precedence analysis, and later during dimension analysis. This information is used later to determine the execution sequence.

Various types of relationships among program entities
have direct implication on the execution sequence of their
corresponding program events. The precedence relationships
among the program events are found based on the analysis of
the program entities. For example, a hierarchical
relationship exists when one data entity contains another,
such as when a file contains a record, a record contains a
field, ..., etc. A dependency relationship exists between a
field and an assertion when the field is either a source
variable of the assertion or its target variable. There are
also relationships between data entities and their
associated control variables. The events and their
precedence relations are represented by a directed graph
called an Array Graph.

The Array Graph is created by two procedures, ENHRREL
and ENEXDP. The ENHRREL routine analyzes data description
statements and finds the precedence relations caused by the
hierarchical relations between data entities. The ENEXDP
routine analyzes assertions and finds the precedence
relations from the dependency relations among data fields
and assertions. It also finds the precedence relations
among data entities and their associated control variables.
Since the Array Graph contains the complete precedence
information, it is used to check the completeness and
consistency of the specification and to determine the
computation sequence.

4.2   REPRESENTATION OF PRECEDENCE RELATIONSHIPS

4.2.1   DICTIONARY

Every program entity has a _full_ _name_ which uniquely
identifies it. Most of the entities have a single component
full name. When two data entities share the same name, it
is necessary to qualify the name with their respective file
names to distinguish them. Two data entities within one
file are not allowed to share the same name. A file name
may have at most two instances denoted as NEW or OLD
followed by an identifier. Thus a data entity may have a
full name of three components: NEW or OLD, file name, and
data name.   Control variables have one component more than
the associated data entities, i.e., a reserved key name.
The full name and the attributes of each program entity are
stored in the dictionary.

In order to use memory efficiently, memory space for
the entries of the dictionary are allocated dynamically.
Pointers to the dictionary entries are stored in a vector
DICTPTR and the total number of pointers in the vector is
denoted as DICTIND. With this arrangement, we can allocate
memory piecewise and access the information randomly. Since
each program entity corresponds to a node in the Array
Graph, we will call its entry number in the dictionary _node_
number. The organization of the dictionary is shown in
Fig. 4.1 and the attributes in the dictionary are listed in

Table 4.1.
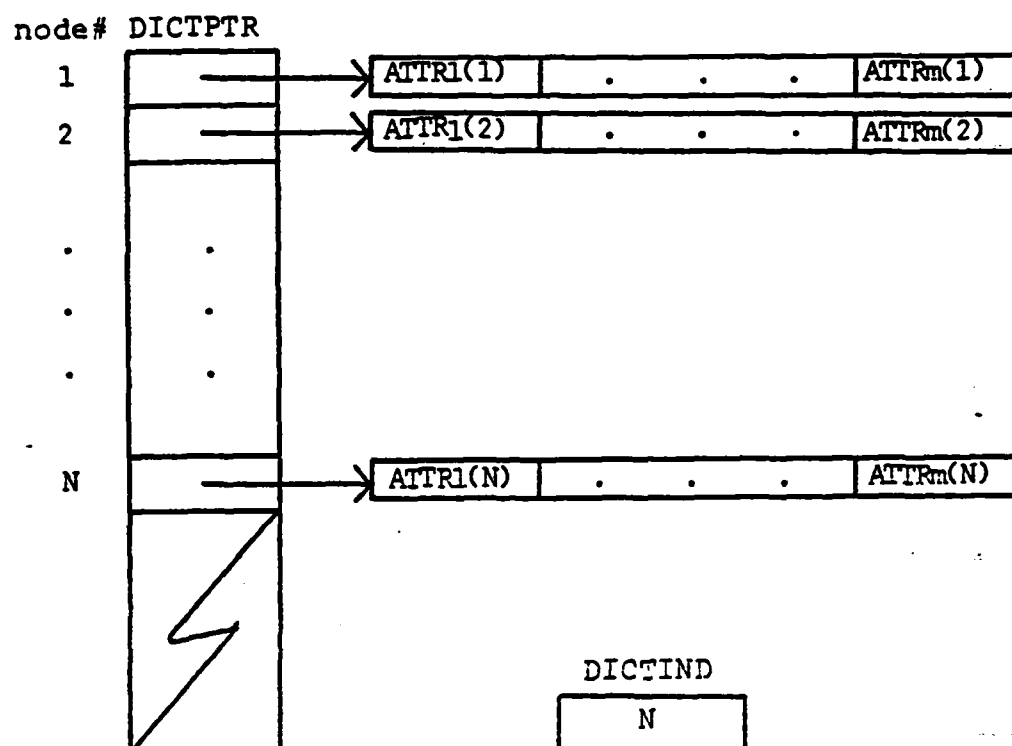
Fig. 4.1 Organization of the dictionary

Table 4.1 Attributes in the Dictionary

XDICT - Is the full name of the entity.

XNAMESIZE - Is the number of characters in XDICT field.

XUNIQUE - Is the smallest name by which the entity can be identified uniquely. If the file name component of a full name is not necessary to identify the entity uniquely, then XUNIQUE is set to the name without file name component; otherwise, XUNIQUE is set to XDICT.

XDICTYPE - Specifies the type of the entity. Following are the possible values:

ASTX - An assertion.

GRP - A group.

FILE - A file.

RECD - A record.

MODL - The specification name.

SPCN - A special name prefixed with a keyword such as END, SIZE, LEN, POINTER, NEXT, SUBSET, ENDFILE, and FOUND.

$SUB - User or system declared subscripts, including the standard subscripts: SUB1, SUB2, ..., SUB10.

$$ - System added subscripts: $1, $2, ..., $10.

$$I - System loop variables: $I1, $I2, ..., $I10.

XMAINASS - Contains a pointer to the storage of the statement which defines the entity.

Table 4.1 Attributes in the Dictionary (Continued)

XNRECS - This count is meaningful only for file entities and holds the number of different record types contained in the file.

XPARFILE - Holds the node number of the parent file entity for all input and output data items.

XPAREC - For data items below the record level this field holds the node number of their parent record entity.

XINP - Is '1'B if the entity is in input file, and '0'B otherwise.

XOUP - Is '1'B if the entity is in output file, and '0'B otherwise.

XISAM - Is '1'B if the entity is an ISAM file, and '0'B otherwise.

XKEYED - Is '1'B if the data entity is in a file for which a key name was specified.

XLEN_DAT - The length in bytes of the data entity.

XREPTNG - Is '1'B if the data entity is repeating.

XVARYREP - Is '1'B if the data entity has a varying number of repetitions.

XMAX_REP - The maximal number of repetitions which was declared for the data entity. If no maximal repetition is declared, XMAX_REP is set to 1.

XVARS - Is '1'B if the entity contains a descendant below the record level and the descendant has a variable structure.

Table 4.1 Attributes in the Dictionary (Continued)

XSUBREC - Is '1'B if the data entity is a member of some record type.

XISSTARRED - Is '1'B if the data entity is repeating and has a undetermined repetition.

XFATHER - The node number of the data entity which is one level above the current entity in the data structure.

XSON1 - The node number of the leftmost descendant of the current entity.

XBROTHER - The node number of the immediate right neighbor of the current entity in the data structure.

XENDB - The node number of the control variable END.X if the currnt entity is X.

XEXISTB - The node number of the control variable SIZE.X if the current entity is X.

XVIR_DIM - The conceptual (virtual) dimensionality of the entity.

XSUBSLST - A pointer to the node subscript list associated with the entity.

X$SUCCESSORS - The number of edges in the XSUCC_LIST.

XSUCC_LIST - A pointer to the list of edges emanating from the current entity.

X$PREDECESSORS - The number of edges in the XPRED_LIST.

XPRED_LIST - A pointer to the list of edges coming into the current entity.

4.2.2 THE ARRAY GRAPH

The Array Graph is a directed graph which represents the precedence relationships among program events. The nodes in the Array Graph are the program events and the edges are the precedence relationships. One program event in the Array Graph will correspond to one program entity. Thus the nodes in the Array Graph correspond to the program entities in the dictionary. The edges between nodes are stored in edge lists associated with those nodes. The attribute SUCC_LIST of a node contains a list of edges emanating from it and the attribute PRED_LIST contains a list of edges terminating at this node. We can thus find the successors as well as the predecessors of any node.

The nodes in the Array Graph are compound nodes, i.e., an entire array of data is represented by one node. Also each assertion is represented by one node, independently of how many array elements it defines. The range of each dimension of a compound node is stored in the node subscript list associated with the node. The edges in the Array Graph are compound edges which denote arrays of relations between two compound nodes. With each edge are also stored the types of subscript expressions used in the relations between the source and the target node of the edge. The meaning of the Array Graph is made more precise by considering the corresponding Underlying Graph (UG), where every array element is represented by one node. An assertion node in

the Array Graph may be expanded in the UG into as many nodes as the elements of the array which it defines. Edges are drawn between the simple nodes. The UG may be an enormous graph which is impractical to analyze. Sometimes the actual number of array elements is not known until run time. Thus it is impossible to create the UG of the specification. In contrast, the Array Graph is more compact and easy to analyze.

4.2.2.1 DATA STRUCTURE OF EDGES

Every edge from a node S to a node T has a uniform format:

$$T(U1, \ldots, Uk) \xleftarrow{t} S(J1, \ldots, Jm)$$

where t is the type of the edge,

  k is the dimensionality of node T,

  m is the dimensionality of node S,

  Ji, 1<=i<=m, are subscript expressions appeared on
      the ith dimension of node S.

  Ui, 1<=i<=k, are the node subscripts associated with
      the node T.

The subscripts U1, ...,Uk of the target node T are stored in the attribute XSUBSLST of T in the dictionary. Therefore they are not specified in the edge. In the later discussion, a type 4 subscript expression Ji will be

indicated by an '*' in the ith dimension of the source node.
An edge is represented by the following data structure:

SOURCE : The source node of the edge.

TARGET : The target node of the edge.

EDGE_TYPE : The type of the edge.

DIMDIF : The difference between the dimensionality of
the target node and the source node.

SUBX : A pointer to the subscript expression list
$(J_1, \ldots, J_m)$.

4.2.2.2 DATA STRUCTURE OF SUBSCRIPT EXPRESSION LIST

A subscript expression J_i can be classified into one of
the following seven categories according to its composition
(refer to section 3.3.2). Type 4 subscript expression is
referenced later as a general subscript expression. Types
5, 6, and 7 subscript expressions are added for the
efficient implementation of some list type
functions[PNPR 80]. They are basically of the form $X(I)$
where X is a variable but used to subscript another variable
B in $B(X(I))$. This form of subscript expression is referred
to as indirect indexing. The array used in indirect
indexing must be integer valued with non-negative entries.
The system will analyze indirect subscripts only if the
indirect indexing array $X(I)$ is sublinear, namely if it is:
a) Monotonic, i.e., if $I>J$ then $X(I) >= X(J)$.

b) Grows more slowly than I, i.e., X(I) <= I.

The system can test the indirect indexing array automatically to determine if it is sublinear by the following simple criteria. In the assertion that define the indirect indexing array X(I), the value of the right hand side must be either 0 or 1 for I=1 and must be equal to X(I-1) or X(I-1)+1 for I>1. Thus the system will examine the assertion to check if it is in the form:

```
X(I) = IF I=1 THEN (1 | 0)
                ELSE (X(I-1) | X(I-1)+1) ;
```

An element in a subscript expression list is defined by the following data structure:

NXT_SUBL : A pointer to the next element of the list.

LOCAL_SUB$: If the subscript expression is of the form Uq[-c] or X(Uq[-c])[-k], then LOCAL_SUB$ is q, i.e. the ordinal number of the subscript Uq as it appears in T(Uk,...,U1).

APR_MODE : The type of subscript expression.

INXVEC : The node number of the indirect indexing vector X if the APR_MODE is 5, 6, or 7. Otherwise, 0.

4.3   CREATION OF THE DICTIONARY (CRDICT)

The procedure CRDICT analyzes the statements of the specification and enters all the program entities into the

dictionary. To find all the data entities we start from the top level of data structures and then trace down the structures. The structures whose root is a file listed in the SOURCE FILE or TARGET FILE statements of the program header are considered external files, i.e. input file or output file. If a data structure is not part of any input or output file, it is considered an interim variable which is computed as any variable in an output file but not written to the external storage.

Corresponding to each input or output file, there is a file entity entered into the dictionary. If a file named F is served both as a source and a target file, then two file entities named OLD.F and NEW.F will be entered into the dictionary. Starting from the file entity we can find its immediate descendants from the file description statement, and the descendants' names will be prefixed by the file entity's name. If the root of a data structure is not a file, we will consider INTERIM as its file name and all the decendants will be put into dictionary, too.

As we analyze a data structure, we also construct a tree representation for it. For every data node we store pointers to its father, leftmost son, and younger (i.e. immediate to its right side) brother in the attributes XFATHER, XSON1, and XBROTHER respectively. We will illustrate this with an example in Fig. 4.2.

```
X IS GROUP (Y,Z) ;

Y IS FIELD ;

Z IS FIELD ;



X = XFATHER(Y)

X = XFATHER(Z)

Y = XSON1(X)

Z = XBROTHER(Y)
```
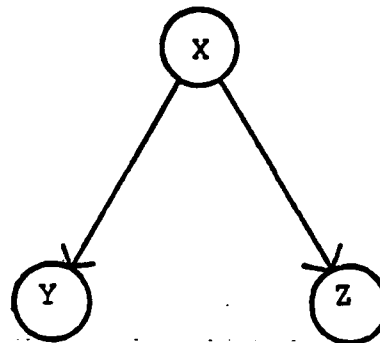
Fig. 4.2 Tree representation of data structure

After all the data entities are entered into the dictionary, a **simplified name** is derived for every data entry. If the file name component can be omitted from the full name without causing any ambiguity, the simplified name is the reduced name. Otherwise the simplified name is the same as the full name.

Other types of program entities such as module name, assertions, and subscript variables are defined by a specific type of statement respectively and there is a one-to-one correspondence between the statements and the entities. We can retrieve these types of statements from the associative memory and enter the entities into the dictionary.

Finally we will put control variables into the dictionary. For each type of qualifier keyword, we find from the program specification all the qualified names with that qualifier. Next we search the dictionary for the suffix name. If the suffix is a declared data entity, the full name of the control variable is formed from the full name of the associated data entity. Otherwise, the qualified name is an unrecognizable symbol and is reported as such to the user.

4.4 CREATION OF ARRAY GRAPH

4.4.1 ENTER HIERARCHICAL RELATIONSHIPS (ENHRREL)

The data stored in external sequential files are simply a string of bits. The use of data description statements allows the user to treat them as structured. Therefore, the system has to transform the data files from a linear form to the structured form which is described by the user. For this purpose, we envisage that there are two program events corresponding to each data entity, one for opening the data and the other for closing the data. The sequential order of data in the external file requires these opening and closing events be arranged in a strict order. The precedence relationship among these program events can be established as follows. If a data entity contains some members, then its opening event precedes the opening event of its first

member and its closing event follows the closing event of
its last member. In addition, the closing event of its nth
member precedes the opening event of its n+1th member. In
the case that a data entity is repeating, then the closing
event of its n-1th instance precedes the opening event of
its nth instance. Fig. 4.3 shows the precedence
relationship of a sequential file. Because the data node B
is repeating, there is an edge from the n-1th instance of
the closing event of node B to the nth instance of the
opening event of node B. The edge is shown as a dashed
line. The existence of this feedback edge causes a cycle in
the Array Graph and this cycle ensures us that the reading
of an instance of the field D will be followed by the
reading of an instance of E. It should be noted that the
subscript expression associated with the edge from the event
C.B to the event O.B is of the form I-1 which allows us to
remove it and break the cycle during the scheduling phase.

```
A IS FILE (B(*),C(*)) ;

B IS RECORD (D,E) ;

C IS RECORD (F,G) ;

D,E,F,G ARE FIELD ;
```
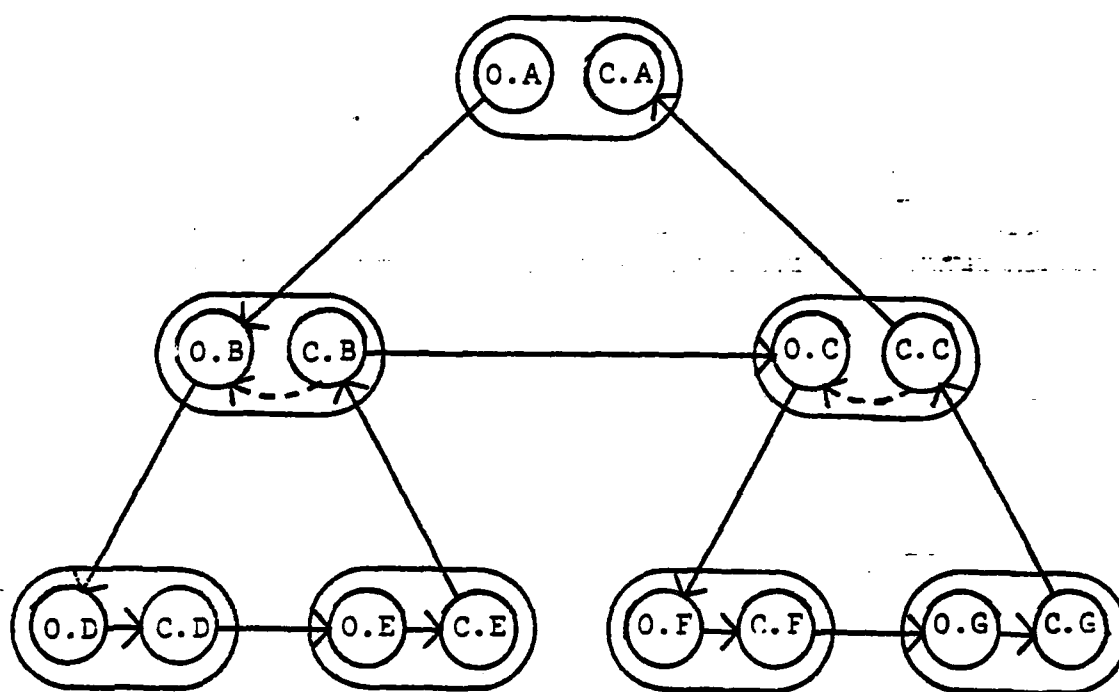

* O.X: opening event for data X

* C.X: closing event for data X

Fig. 4.3 Precedence relationship of a data structure

We envisage that for each field entity there is a third node which corresponds to the available event of the data. The opening event of an input field must precede its available event, and the closing event of an output field

should follow its available event.

This view assures us that we can always read the input files sequentially and store them in the main memory before any computation starts. If there are variable structures, i.e., structures of varying field length or varying number of repetitions, then we may have to include some assertions in the reading process. Afterwards we can do all the computation internally conforming with the constraint of data dependency which is implied by the assertions. At the end, all the fields in the output files are available and the informations for controlling the variable structure are available, too. We then take the data from main memory, assemble them into records, and write the records sequentially.

Actually we have in the Array Graph only one node, instead of the open, close, and available nodes mentioned above, for each data entity, as this helps compiler efficiency. For input files, we can view the nodes as corresponding to the opening events. For output files, the nodes corresponding to the closing events. The records stored in a sequential file have to be accessed in a strict order. Therefore, there is a precedence relationships among the data entities of an input or output file to assure that the records are accessed in the proper order. On the other hand, a record is composed of fields. The __membership relation__ between a record and its constituent fields implies

a precedence relationship, i.e. no field in an input record will be available until the record is read in. Similarly all the fields in an output record should be available before the record can be written out.

We will use the following definitions in discussing tree structures.

Definition For a data entity G, SON1(G) denotes its leftmost son.

Definition For a data entity G, RSON(G) denotes its rightmost son.

Definition For a data entity G, CEB(G) denotes the closest elder brother of G, i.e. the data entity which is to the immediate left of G among all the brothers of G.

Definition For a data entity G, CYB(G) denotes its closest younger brother, i.e. the data entity which is to the immediate right of G among all the brothers of G.

Definition For any tree with node G as the root, RDM(G) denotes the rightmost node on the frontier of the tree.

Definition For any tree with node G as the root, LDM(G) denotes the leftmost node on the frontier of the tree.

The precedence relationships in different file types is discussed in the following.

1) Input sequential file. Since the records in a sequential file are read in one at a time, the precedence relationship needs to assure that the records are read in the order they are present in the input file. A record may be composed of many fields. Therefore, after a record is read, it should be unpacked to get all the fields. If the records in a file are not unpacked in the order they are read, then we will need memory space to store the records. Therefore, it is advantageous to unpack the records when they are read in. This implies that all the fields in a sequential file will become available in the order they occur in the external file. Three kind of edges are drawn among the data nodes in an input sequential file.

a) Assume that a data node G is n dimensional. If SON1(G) exists and is m dimensional where m may be either n or n+1, then the following edge is drawn.

$$SON1(G)(J1,\ldots,Jm) \; \text{<-1a-} \; G(J1,\ldots,Jn)$$

b) Assume that a data node G is n dimensional and FATHER(G) is k dimensional where k may be either n-1 or n depending on whether node G repeats or not. If CEB(G) exists and RDM(CEB(G)) is m dimensional, then the following edge is drawn.

$$G(J1,\ldots,Jn) \; \text{<-1b-} \; RDM(CEB(G))(J1,\ldots,Jk,*,\ldots,*)$$

c) Assuming that a data node G is n dimensional. If it is repeating, then the following edge is drawn.

$$G(J1,\ldots,J_n) \; \text{<-1c-} \; RDM(G)(J1,\ldots,J_n-1,*,\ldots,*)$$

If a data node in an input sequential file corresponds to the opening event of that data, we can interpret the above edges in the following way. The edges of type 1a say that a higher level data instance should be ready before all of the data instances corresponding to the first member of it can be read. The edges of type 1b say that all the brothers within the same instance of their father should be read in the order they are declared in the data structure. The edges of type 1c say that if a data node is repeating, then one instance of it is not ready to be read until the last field in the previous instance of it is read.

2) Output sequential file. The records of an output sequential file should be written out in a strict order. There may be several fields in a record, therefore, we may have to pack the fields before writing. Packing the fields when they become available is convenient for the code generation but poses extra restrictions on scheduling the assertions. For example, suppose a record node R contains three fields A, B, and C. If we insist that fields A, B, and C should be available in that order, the user would not be able to define the value of A in terms of C. Therefore, at or above the record level the precedence relationship requires that the records be written in strict order but below record level the

AD
A111809

1.0

1.1

1.25   1.4   1.6

2.8   2.5

3.2   2.2

3.6

4.0   2.0

1.8

MICROCOPY RESOLUTION TEST CHART
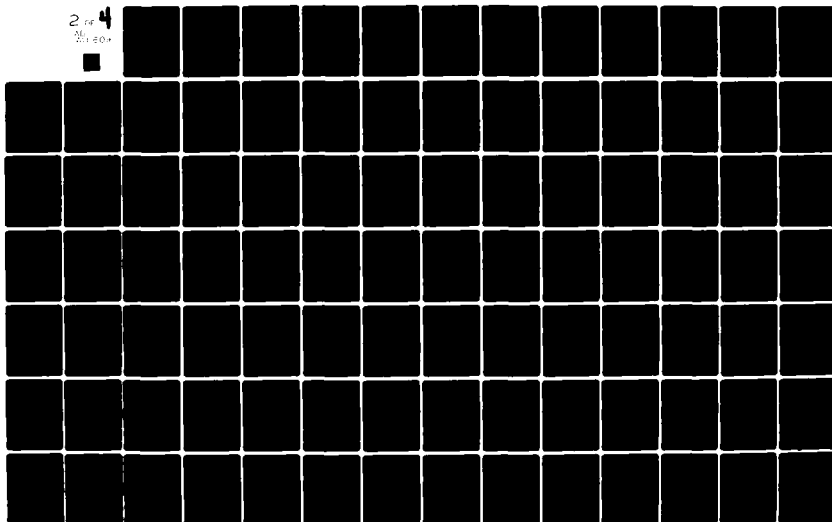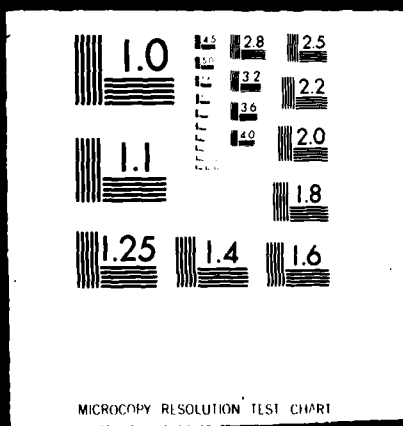
precedence relationship will only require that the constituent fields of a record are ready before the record is written. Therefore, fields in a record do not have to be computed in the order they are packed into the record.

Three kinds of edges are drawn among the data entities above and including the record level of an output sequential file.

a) Assuming that G is an n dimensional data entity above the record level and RSON(G) , i.e. the rightmost son of G, is m dimensional. The following edge is drawn from RSON(G) to G.

   G(J1,...,Jn) <-2a- RSON(G)(J1,...,Jn,*)

b) If node G has a younger brother, then an edge will be drawn from node G to LDM(CYB(G)). Let G be an n dimensional node, FATHER(G) be a k dimensional node, and LDM(CYB(G)) be a m dimensional node. The edge to be drawn is as follows.

   LDM(CYB(G))(J1,...,Jk,...,Jm) <-2b- G(J1,...,Jk,*)

c) If node G is repeating, then the following edge is drawn from G to LDM(G). Let G be an n dimensional node and LDM(G) be a m dimensional node.

   LDM(G)(J1,...,Jn,...Jm) <-2c- G(J1,...,Jn-1)

If we imagine that a data node in an output sequential file corresponds to the closing event of that data, then the edges mentioned above have the following

interpretation. An edge of type 2a says that a data instance can be written out only after all the data instances corresponding to its last son are written out. An edge of type 2b says that all the instances of an elder brother within the same father instance should be written before any instance of its younger brother can be written. An edge of type 2c says that if a data node is repeating, then an instance of it cannot begin to be written until the previous instance is completely written.

Below the record level in an output file, the precedence relationships assures that a record will not be written out until all of its constituent fields are available. However, the relative order in which the fields are computed is not restricted. We will simply draw edges from all the descendants of a record node to it. Fig. 4.4 illustrate the edges in an output sequential file.

```
A IS FILE (B(*),C(*)) ;

B IS RECORD (D,E) ;

C IS RECORD (F,G) ;

D,E,F,G ARE FIELD ;
```
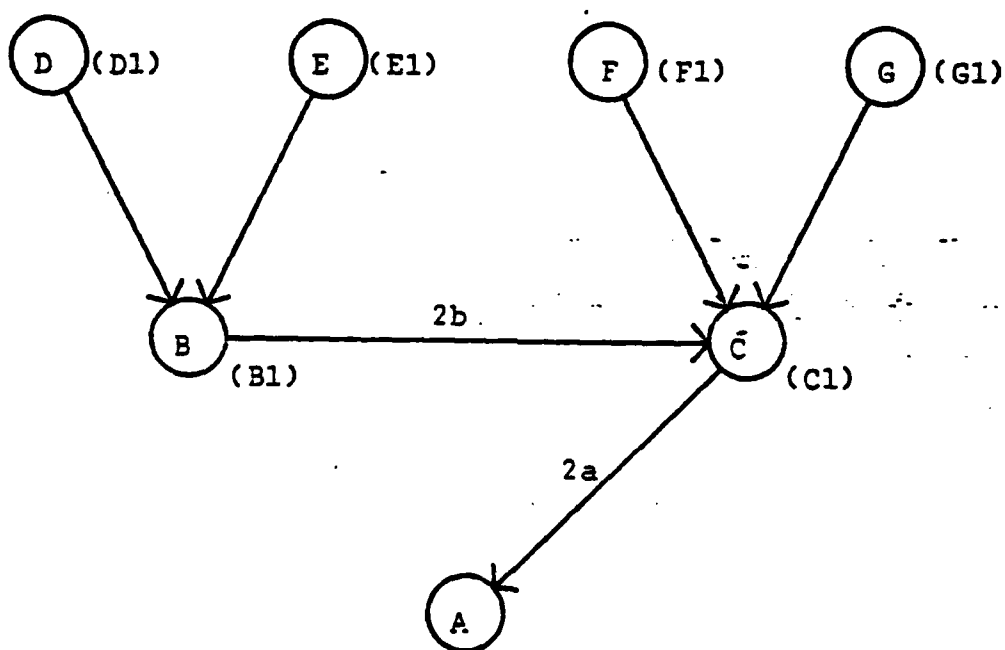

Fig. 4.4 The edges in an output sequential file

3) An input ISAM file. In an ISAM file, there is only one
type of record. The dimensionality of the record node IR
is the same as that of the associated control variable
POINTER.IR. Since the record instances are accessed with
the keys, it is possible to read the records in the order
of the keys. If the ISAM file is a pure source file to
the program, the keys in the POINTER.IR array can be used
in any order. On the other hand, if the ISAM file is

used as a source and target file, the records should be processed in a sequential way, therefore, the keys in the POINTER array should be used sequentially to access the records. Below the record level, we can have the similar precedence relationship as in a SAM file because we may have to unpack the fields.

4) An output ISAM file. If an ISAM file is a pure target file, the output records will be added to the file. If it is a source and target file to the program, then only the selected records may be updated. In order to assure that each updated record includes the effects of previous updates, we will have to update and write out a record before the next record is read in. Therefore, the keys in the POINTER array should be used sequentially. However the fields in an output record can be computed in any order. Below record level the precedence relationships only reflect the membership of the fields within the record.

5) Interim variable. There are no I/O actions concerning interim variables. They are stored in main memory and referenced as fields. Therefore, there is no relative precedence relationship among the interim fields. But we still draw edges which reflect the membership among the data entities to facilitate range propagation (refer to Chapter 5). Since an interim variable is considered to be part of an output file except that it will not be

written out, the edges are drawn from the descendants to
the ancestors.

4.4.2 ENTER DEPENDENCY RELATIONSHIPS (ENEXDP)

Two types of assertions, namely simple assertion and
conditional assertion, may be used to define the values of
interim variables and output variables. The execution of an
assertion depends on the availability of all of its source
variables, and its execution makes the target variable
available.   This is because a data entity must be defined
before it is referenced and a data entity becomes available
after the assertion in which it is the target variable is
executed.

Procedure ENEXDP examines all the assertions twice. In
the first pass, it checks whether the target variable of an
assertion defines a sublinear function and can be used as an
indirect indexing vector or not. An indirect indexing array
should be defined by an assertion of the following form.

X(I) = IF I=1 THEN (0 | 1)

ELSE (X(I-1) | X(I-1)+1) ;

During the second pass, it analyzes every assertion and
enters the precedence relations caused by explicit data
dependency into the Array Graph. Given a simple assertion,
the left hand side of it is scanned to find the target

variable. Then the expression on the right hand side is scanned to find all the source variables. For a conditional assertion, the THEN parts, ELSE parts, and the conditional expression parts are scanned in that order to find all the source and the target variables. The source variables in a conditional assertion are found in the conditional expressions, the THEN parts, and the ELSE parts. For every source variable an edge is drawn from it to the assertion node. It should be noted that one assertion defines one target variable only and no more than one target variable can appear in a conditional assertion.

The edge from the source variable to the assertion is of EDGE_TYPE 3 and the edge from the assertion to the target variable is of EDGE_TYPE 7. The DIMDIF is the dimensionality difference of the target node and the source node of the edge. The types of the subscript expressions of a source variable are stored in the subscript expression list associated with the edge. It should be noted that the subscript expressions of the target variable define a mapping from the node subscripts of the target variable to the node subscripts of the assertion. Because the edge corresponding to the occurrence of the target variable is drawn from the assertion node to the target variable, instead of from the target variable to the assertion node, the mapping should be inverted to form the subscript expression list of the edge. In Fig. 4.5 the data

dependency of an assertion is shown. Notice that there is a list of subscripts associated with every node in the graph. For example, variable A is a two dimensional array. Subscripts $\langle A,1 \rangle$ and $\langle A,2 \rangle$ correspond to the first and second dimension of array A. The edge leading from node A to al has a subscript expression list associated with it. The subscript expressions are ordered in the way they are used in the subscript variable $A(I,J-1)$.

al: $C(I,J) = A(I,J-1) + B(I,4)$;
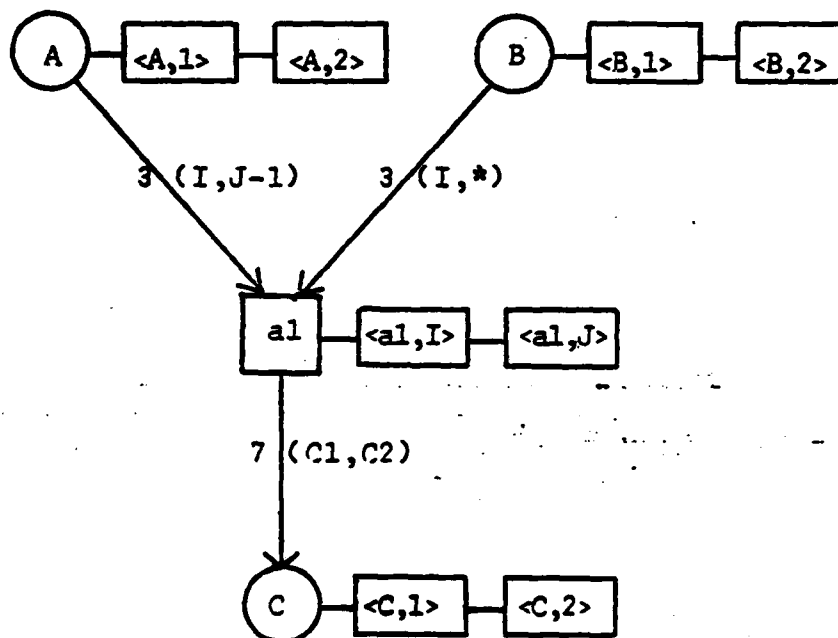
Fig. 4.5 The data dependency of an assertion

In addition to the explicit data dependency found in an assertion, there exists some implicit data dependency between the data entities and their associated control

variables. Let TRGT denote the name of a data entity and NODE denote the name of the associated control variable which is composed of a keyword PREFIX followed by the name of the data entity.

1. If PREFIX = 'POINTER', then verify that TRGT is a keyed record and draw an edge.

    TRGT <-5- POINTER.TRGT,   DIMDIF = 0 .

2. If PREFIX = 'SIZE', then verify that TRGT is repeating and draw an edge.

    TRGT(I) <-13- SIZE.TRGT,   DIMDIF = 1 .

3. If PREFIX = 'END', then verify that TRGT is repeating and draw an edge.

    TRGT(I) <-14- END.TRGT(I-1), DIMDIF = 0 .

4. If PREFIX = 'FOUND', then varify that TRGT is a keyed record and draw an edge.

    FOUND.TRGT <-15- TRGT, DIMDIF = 0 .

5. If PREFIX = 'NEXT', then verify that TRGT is a field in an input sequential file and draw an edge.

    NEXT.TRGT <-16- TRGT, DIMDIF = 0 .

6. If PREFIX = 'SUBSET', then verify that TRGT is an output record. If it is an output record, then draw the following edge.

    TRGT <-17- SUBSET.TRGT, DIMDIF = 0 .

7. If PREFIX = 'LEN', then we draw an edge.

    TRGT <-20- LEN.TRGT, DIMDIF = 0 .

The subscript expression lists of these edges are for the moment empty. They will be constructed by the procedure FILLSUB later according to the EDGE_TYPE.

4.5 FINDING IMPLICIT PREDECESSORS (ENIMDP)

Many efforts have been made to make MODEL language tolerate some incompletenesses and inconsistencies in the specification. When incompletenesses and inconsistencies are found, warning messages or error messages are sent to the user. If practical, the MODEL processor tries to correct the specification in a reasonable way.

If an interim field is not defined by any assertion, an error message is sent to inform the user. It is probable that the user forgot to write the assertion. Therefore, the system should request an assertion from the user. However, if a field in a target file is not defined explicitly, the MODEL processor will try to find an implicit source to define that field. The MODEL processor tolerates this kind of incompleteness and saves the user work of writing assertions for merely copying fields from a source file to a target file.

Given a field in a target file which is not explicitly defined by any assertion, we will search for a field with the same name in another file according to the following

order of priority. The idea is to make some reasonable assumption so that the undefined field will get a value.

Rule 1: If the undefined field is in a file which is both a source and target file, then the value in the corresponding field in the old record is taken as the value for it.

Rule 2: If Rule 1 does not apply, then the processor tries to find a same-named field in other source files. If one is found, it is assumed to be the source. If more than one is found, then the processor arbitrarily picks one as the source and prints a message to indicate that there was ambiguity.

Rule 3: If the above are unsuccessful, the processor tries to find a field with the same name in other output files. If one is found, it is taken as the source, and if more than one is found, then one is taken arbitrarily, with a corresponding message to the user regarding the ambiguity.

In the above cases where an implicit predecessor is found successfully, an assertion which defines the target variable by the implicit predecessor is generated as if it were entered by the user.

4.6 DIMENSION PROPAGATION (DIMPROP)

The source and the target variables in an assertion may be arrays. In order to reference an element of an N dimensional array, the user should subscript the array name with N subscript expressions. A subscriptless dialect of the MODEL language allows the user to omit subscripts in assertions in certain cases which do not lead to ambiguity. Therefore, the number of subscript expressions following an array variable does not necessarily indicate its actual dimensionality. Furthermore, the declaration of a multi-dimensional interim array may be simplified by omitting the data description statements for the higher level groups. The omission of subscript expressions in assertions and the omission of the higher level data description can be viewed as incompleteness or inconsistency of the specification. However, they are tolerated by the MODEL processor, and a process called dimension propagation is used to resolve inconsistencies of the dimensionality for the interim variables and missing subscripts in assertions.

All the nodes in input and output files should be declared precisely, using data description statements. Their number of dimensions can therefore be derived directly from the data description statements. Associated with every edge there is a field DIMDIF which denotes the dimension difference between the source and the target nodes of the edge. The number of dimensions of a node can be propagated

along the edges of the Array Graph.

The dimension propagation algorithm is briefly described in the following. Let N denote the set of nodes in the Array Graph, array C store the current number of dimensions, and array D store the initially declared number of dimensions for each node in N. A queue Q keeps all the nodes whose calculated dimension could possibly be changed.

Algorithm 4.1 Dimension Propagation

Input. Array Graph.

Output. VIR_DIM: An attribute in the dictionary which contains the number of dimensions of a node.

1. For each node n in N, let C(n) be D(n) and put node n in Q.

2. If Q is empty, then exit.

3. Pick a node n from Q, remove it from Q. Let dim be 0.

4. For every incoming edge from node s to n, let dim be the maximum of dim and C(s)+DIMDIF.

5. For every outgoing edge from node n to t, let dim be the maximum of dim and C(t)-DIMDIF.

6. If dim<=C(n), go to step 2.

7. Else, the node n has a new updated dimension. Let C(n) be dim.

8. For every incoming edge from node s to n, append s to Q.

9. For every outgoing edge from node n to t, append t to Q.

10. If more than N*N nodes have been taken from the queue, then halt and issue an error message - there exists a

propagation cycle.

If the process converges, then every node will have a finite dimension. However, it is possible that a cycle in the graph causes an endless increase in the dimensions. Consider for example the following specification.

(F, H) ARE FIELD ;

I IS SUBSCRIPT ;

IF I=1 THEN H(I) = 5 ; ELSE H(I) = F+1 ;

IF I=1 THEN F(I) = 6 ; ELSE F(I) = H+1 ;

The first assertion implies that the dimension of H is larger by 1 than that of F, i.e. C(H)>C(F). The second assertion states that C(F)>C(H). Applying our algorithm to this specification will result in endless loop of alternately incrementing C(H) and C(F). In this case the system will send out an error message indicating that the dimension propagation process is in an infinite cycle and also print out the nodes involved in the cycle.

4.7 FILLING MISSING SUBSCRIPTS IN ASSERTIONS (FILLSUB)

In the dimension propagation phase we have determined the number of dimensions of every node. If the number of dimensions of a node is larger than its apparent number of

dimensions, it is necessary to add the respective subscript and data structures. This is performed in the following three tasks.

Task 1: Generate the node subscript list.

If the node X is a data node, its node subscript list is (displayed here from last to first):

(FOR_EACH.Ak, ... , FOR_EACH.A1)

where Ak, ..., A1 is the list of the repeating ancestors of X in a top down order. If X itself is repeating than A1 is equal to X.

If the node is an assertion node, then it has already been assigned a partial subscript list by ENEXDP. This is the list of apparent subscripts in the assertion, i.e. all the subscripts appearing either on the L.H.S. or the R.H.S. of the assertion. Let the assertion be of the form:

a1: A(Ik, ..., I1) = f(....) ;

Let the R.H.S. contains the subscripts J1, ..., Jm not appearing on the L.H.S. and hence assumed to be reduced. Then the partial list assigned to a1 is (Ik, ..., I1,Jm, ...,J1) and its apparent dimensionality is determined to be $d=k+m$. As a result of the dimension propagation process we may have recomputed a new dimensionality c for a1 where $c>=d$. This will cause $n=c-d$ new subscripts to be added to the subscript list of a1 which now appears as:

($n, ..., $1,Ik,...I1,Jm,....,J1)

where $1, ..., $n are the name of the new subscripts.

Task 2: Fill in Missing Subscripts in the Assertions.

Consider an instance of a subscripted variable A(Ij, ..., Il) in an assertion. The calculated dimension VIR_DIM for array A yields a value d which should be greater or equal to j. If n=d-j>0 we should add n new system generated subscripts $1 to $n, modifying the instance into A($n, ..., $1,Ij, ..., Il). It should be noted that the new subscripts are always added on the leftmost dimensions of the array variables.

Task 3: Fill in the Subscript Expression List for the Edges.

All the edges except types 3 and 7 have been generated with an empty subscript expression list. Using the edge type and the dimensions of its source and target nodes, we generate a subscript expression list for each edge. Edges of type 3 and 7 have a partial subscript expression list based on their apparent appearance in the assertion. It may be necessary to expand this partial list. If n missing subscripts have been added to the variables in an assertion, then it is necessary to add n subscript expressions to the edges which correspond to the instances of the variables in the assertion.

CHAPTER 5

RANGE PROPAGATION

5.1   INTRODUCTION

The structures of variables are declared in data
description statements.   Every variable is considered an
array of some dimensions. The number of elements in an
array variable is determined by the dimensionality of the
array and the sizes of each of the array dimensions.   The
size of an array dimension is called the <u>range</u> of that
dimension. The range information allows us to allocate
memory space for the array variables and generate iteration
control statements which will define every element in the
arrays.   The use of subscripts in assertions makes it
possible to define multiple elements of an array through one
assertion.   We can instantiate an assertion by fixing its
subscript values. Then every instance of the assertion
defines one single data element.   The ranges of the
assertion's subscripts restrict the number of instances of
an assertion, which in turn defines the number of times that

the assertion will be executed. The ranges of array dimensions and assertion subscripts are used in the later phases to synthesize the program.

Much information is not given explicitly in the specification. For instance users are allowed in assertions to use free subscripts for which the range is not specified. Also the range specifications of some array dimensions may be omitted. Therefore an algorithm is needed to derive ranges for certain assertion subscripts and array dimensions.

There is yet another reason why we want to analyze the subscript ranges. A criterion for placing a number of assertions in the scope of one loop is that they all have subscripts of the same range. From the point of view of program optimization it is preferred to have the loop scope as large as possible. It is important therefore to identify the subscripts of the same range. By propagating the specified range information to all the assertion subscripts and array dimensions we not only find the ranges which have been incompletely specified, but also identify the ranges which are equal.

5.2 LANGUAGE CONSTRUCTS FOR RANGE SPECIFICATION

A multi-dimensional array is declared as a hierarchical data structure with the most significant dimension specified at the top level. The range of a dimension may not depend on the subscript value of less significant dimension. The range of an array dimension may be specified in MODEL in several alternate ways as follows:

(1) Through a data description statement. A constant number of repetitions of a data structure may be specified in the data description statement which describes the parent structure.

(2) By defining the value of a SIZE qualified control variable (Refer to section 3.4.). For example, if group X repeats M times and M is a variable itself, we may use the following assertion to specify its range:

SIZE.X = M ;

A SIZE qualified variable is an interim variable of at most one dimension less than that of the suffix variable. Its value is used to define the range of the last dimension of the suffix variable (i.e. X). Consider an N dimensional repeating group X. Assume that the ranges of all its dimensions except the least significant one are defined elsewhere. By definition, SIZE.X is at most an N-1 dimensional array and the range of its dimensions is exactly the same as the range of corresponding dimensions of data structure X. Since the

values in array SIZE.X can be different from one another, the array X may not have a regular (i.e. rectangular) shape, but have "jagged edges." This can be stated formally as follows:

$$X(S_1, S_2, \ldots, S_k, \ldots, S_n) \text{ is in X iff}$$

$$SIZE.X(S_1, \ldots, S_k) \text{ is in SIZE.X } \&$$

$$1 <= S_n <= SIZE.X(S_1, \ldots, S_k)$$

(3) By defining the value of an END qualified control variable. The END array is of boolean type. It determines the range of the least significant dimension of the variable named in the suffix. Given an N dimensional array X, the associated control array END.X has the same structure as array X. The range of the Nth dimension is defined as the smallest positive integer Ln which satisfies the following conditions.

$$END.X(S_1, \ldots, S_{n-1}, Ln) = TRUE \ \&$$

$$END.X(S_1, \ldots, S_{n-1}, S_n) = FALSE,$$

$$\text{for } 1 <= S_n < Ln.$$

(4) By using a subscript declaration statement to define a global subscript. The constant number of repetition can be specified in the statement. For example:

    I IS SUBSCRIPT (20) ;

(5) By system default. A repeating data structure which is a rightmost decendant and which is above or at the record level, may be assigned the end-of-file as its range if the user does not specify a range for it.

The mechanisms of SIZE and END arrays are not totally redundant. There are some essential differences between the SIZE and END arrays. First, the END array can define a minimum range of one, whereas the SIZE can define a range of zero. This is because the END array must have at least one value of boolean true. Secondly, the range specified by SIZE array is finite. But the range specified by END array may be infinite (through a user error in the range defining assertion, when there is no first boolean true condition). This is not checked by the system. Thirdly, the range specified by array $SIZE.X(I1,..,Ik)$ may not depend on the array element $X(I1,..,In)$, while $END.X(I1,...,In)$ may depend on $X(I1,...,In)$. For example, let $X(1),...,X(k)$ be all the instances of an one dimensional array X whose range is specified by $SIZE.X=k$. In the program, the value of SIZE.X, i.e. k, must be computed before we compute any of the elements of X. If END control array is used, the range is specified by $END.X(1), ... , END.X(k)$, and we only have to ensure that $END.X(I-1)$ is computed before $X(I)$ for $1<I<=k$.

5.3 DEFINITIONS

Subscript variables belong to a special class of variables. While an ordinary variable can assume only a unique value, a subscript variable can take on a range of positive integer values. Subscript variables can be used as indices in array element references or in the same way as ordinary variables to compose complicated expressions. The meaning of subscripts is the same as their meaning in mathematical usage.

The following definitions are used in discussing subscripts.

Definition Let X be an N dimensional array represented in the Array Graph by a node. Let i be a positive integer. The tuple <X,i> is referred to as a <u>node subscript</u>. It denotes the ith dimension of the node of array X. Let al be an assertion node, and I a subscript variable referenced in the assertion al. The tuple <al,I> is referred to as a <u>node subscript</u> for I associated with the assertion node al. If <n,d> is a node subscript, then R(<n,d>) denotes its range.

Node subscripts are grouped into <u>range sets</u>. Every range set contains the node subscripts which have the same range. However no two dimensions of the same node can be

put into one range set even if they have the same ranges
because every range set will later correspond to a level of
nested loops in the generated program and no two dimensions
of the same node can correspond to the same level of nesting
loops.

<u>Definition</u> The range of a subscript that has been declared
as a <u>global</u> <u>subscript</u> is the same in all assertions
where it is used. There can only be one range
associated with a global subscript.

<u>Definition</u> The range of a subscript that has not been
declared as global is fixed within the scope of the
assertion where it is used. It will be called a <u>local</u>
<u>subscript</u>. A symbol used as a local subscript can have
different ranges in different assertions.

There are two types of global subscripts in MODEL. One
is specified by use of the qualifying keyword FOR_EACH in
the prefix and a repeating data structure name in the
suffix. The other is explicitly declared in a subscript
declaration statement. (Refer to section 3.3.2.) The
FOR_EACH type global subscript always has the range of the
repeating data group named in the suffix associated with it.
A user declared global subscript can have its range
specified in the subscript declaration statement. By using
global subscripts in assertions, the user can specify

explicitly the range of assertion subscripts.

Local subscripts are all of the form SUBn where n is a
positive integer. Users do not have to declare local
subscripts (in subscript statement). The use of local
subscripts in an assertion is like that of formal parameters
in a function definition. They can be chosen arbitrarily
within the scope of an assertion. This gives the user
freedom to reuse the subscript names in different
assertions.

5.4   DISCUSSION OF RANGE PROPAGATION

5.4.1   CRITERIA FOR RANGE PROPAGATION

In this section we discuss the conditions for
propagating the range of a subscript from one node to
another. A node subscript refers to either an array
dimension or an assertion subscript. If two node subscripts
are related through some dependency relation and one of them
does not have an explicit range specification, we propagate
the range from one to the other.

Let us consider first a simple assertion :
$B(I) = A(I)$. Three entities are involved : the source
variable A, the target variable B, and the assertion itself.
All of them are one dimensional objects. The assertion
states that the kth instance of the assertion corresponds to

the kth instance of array B for all k in the range of B's dimension. There is a bijective mapping between the instances of the assertion and the instances of the array B. It is therefore very natural to believe that the range of the target variable B is the same as the range of the assertion. Additionally, from the subscript expression I in the term A(I) we can derive that the range of the assertion can be taken from the range of the array A. In short, whenever a simple subscript variable is used as a subscript expression it strongly suggests that we may propagate the range from one node subscript to another.

When a subscript expression of the form I-k is used in an assertion, where I is a subscript variable and k is a positive integer, there exists a one-to-one mapping between values of certain elements indexed by I and I-k. The mapping may be interpreted in two possible ways : assume the ranges of the arrays indexed with I and I-k subscripts are the same, or assume that the variable with the I-k subscript expression has k instances fewer than the variable with I subscript. We have decided to adopt the simpler assumption, that is, the ranges are the same. Therefore we will propagate ranges between the node subscripts indexed by subscript expression I and I-k.

It should be noted that we do not intend to modify or ignore a user specified range of a node subscript. The analysis mentioned above is used for two purposes. One is

to derive a range for a node subscript which does not have
an explicitly specified range. Second is to determine if it
is possible to put two node subscripts into the same range
set when both of them have user specified ranges and the
ranges are the same. When two node subscripts have user
specified ranges, we are interested in finding out whether
their ranges are equal. Since there is no simple way to
determine if two functions are equal in general, we will
only check the assertions which define the range arrays by
the other range array.

5.4.2 PRIORITY OF RANGE PROPAGATION

User specified ranges are associated with repeating
data structures or declared global subscripts. The range
specified for a data node is interpreted as the range of its
least significant dimension. Ranges of node subscripts can
be propagated along a path in the Array Graph from one node
to another based on the following relations between
respective node subscripts.

1. The two node subscripts are both global subscripts and
   have the same global subscript name.

2. One of the node subscripts corresponds to a dimension of
   a data node and the other corresponds to the same
   dimension number of the associated control variable.

3. The two node subscripts occur on the corresponding

dimensions of two data nodes in the same data structure.

4. One node subscript is associated with an assertion node and the other is associated with a source variable of the assertion.

5. One node subscript is associated with an assertion node and the other is associated with the target variable of the assertion.

There may be several alternative paths (and directions) for propagating a range, and the range derived for a node subscript may depend on the choice of a path. The choice of path may also affect the efficiency of the generated program. Therefore, we will propagate ranges according to a priority order which attempts to obtain the highest efficiency. The priority order is as follows.

When a global subscript is used in several assertions, the ranges of the respective node subscripts (in these assertions) are the same. We may consider all the node subscripts with the same global subscript name as a group. Whenever any element in the group has its range defined, we will propagate the range to other elements in the same group. This type of propagation will have the top priority.

Next consider the data nodes and their associated control variables such as SIZE.X, END.X, POINTER.X, LEN.X, ..., etc. The dimensions of the control variables correspond to the dimensions of the variable named in the

suffix from left to right. The corresponding dimensions of a data node and its associated control variables should have the same range. Similarly the corresponding dimensions of a data node and its higher level nodes in a data structure should have the same range.

If the range specification of local subscripts in assertions or array dimensions are not given explicitly, we will derive them by analyzing the respective subscript expressions in assertions. It is preferable to propagate the range from a target variable to an assertion rather than to propagate the range from a source variable to an assertion. Therefore, the range propagation between an assertion node and its target node or between a data node and its associated control variable will have the second priority.

Globally it is preferred to propagate the range from a variable in an output file backward to a variable in an input file than reversely. Thus we will assign the third priority to the propagation from an assertion node backward to its source variables and the fourth priority to the propagation from a data node forward to an assertion node in which it is referenced as a source variable.

Example Let array A be an input file with 20 elements, array C an output file with 10 elements and array B one dimensional interim array. The assertions

al: $B(I) = A(I)$;

```
    a2:   C(I) = B(I) ;
```

may lead us to assign either 20 or 10 as the range for
array B, depending on the point of view taken. As far
as the correctness is concerned, it does not make any
difference whether 20 or 10 is used as the range of
array B. But a smaller range would mean potentially
less memory space and less computation time. Therefore
the latter is more desirable. The range may be
evaluated as follows. Since no global subscripts are
used here, no propagation corresponding to the top
priority can be achieved. The propagation from an
assertion node to the target variable is second
priority, therefore, the range of $\langle C,1 \rangle$ and $\langle B,1 \rangle$ should
be propagated to $\langle a2,I \rangle$ and $\langle a1,I \rangle$ respectively. The
range of subscript $\langle B,1 \rangle$ will be that of $\langle A,1 \rangle$ or $\langle C,1 \rangle$
depends on whether we give higher priority to the
propagation from $\langle A,1 \rangle$ to $\langle a1,I \rangle$ or from $\langle a2,I \rangle$ to
$\langle B,1 \rangle$. Since the latter has the higher priority, the
range is propagated from array C all the way back to the
assertion node a1. (Refer to Fig. 5.1.)

```
a1: B(I) = A(I) ;

a2: C(I) = B(I) ;
```
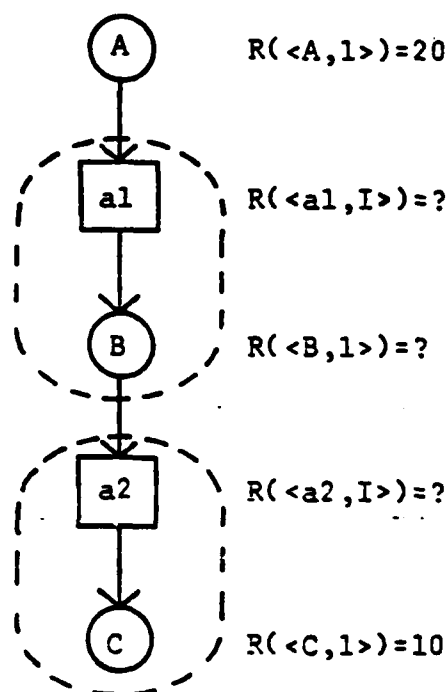

Fig. 5.1 Example of Range Propagation

In summary, we have divided the range propagation into four priority levels. The top level is based on use of global subscripts. The second level is based on the relation between data node and its associated control variables or between the assertions and their target variables. The third level is to propagate the range from an assertion backward to its source variables, and the fourth one is to propagate the range from a data array forward to the assertions in which it is referenced as a source variable.

5.4.3 REAL ARGUMENTS OF RANGE FUNCTIONS

Every node subscript will iterate over its range by a loop control statement in the generated program. A node in the Array Graph having N node subscripts associated with it will have an N level nested loop enclosing it. Every loop controls the iteration of a corresponding node subscript. We will show that the range specification of the node subscripts may have influence on the order that the loops can be nested and on the order of subscripts in referring to a range array.

When the ranges of the dimensions of an array are all constant, the array has a regular shape. We can access all of the array elements by iterating the subscripts in any order. For example, if we have a rectangular array A, we can access all of the array elements either row-wise or column-wise. However, if some of the dimension ranges of an array are specified by range arrays, it is no longer true that we can nest the loops in any order. In Fig. 5.2(a) two arrays A and B are both three dimensional arrays. The ranges of the third dimension of both arrays are specified by the SIZE.A array. In Fig. 5.2(b), a part of the flowchart for the specification in 5.2(a) is shown. The point is that the loop corresponds to node subscript <A,3> should be scheduled inside the loops of <A,1> and <A,2>. Because the loop control statement for <A,3> references the range array SIZE.A and the value of SIZE.A depends on the

values of subscript <A,1> and <A,2>.

```
A IS FIELD;
B IS FIELD;
B(I,J,K) = A(I,J,K) ;
SIZE.A(I,J) = f(I,J) ;
```

Fig. 5.2(a) A range array with real arguments

```
    .
    .
    .
DO <A,1>;
   DO <A,2>;
      DO <A,3> = 1 TO SIZE.A(<A,1>,<A,2>);
         A(<A,1>,<A,2>,<A,3>);
         B(<A,1>,<A,2>,<A,3>) = A(<A,1>,<A,2>,<A,3>);
         B(<A,1>,<A,2>,<A,3>);
      END;
   END;
END;
   .
   .
   .
```

Fig. 5.2(b) Flowchart of 5.2(a)

A simple solution would be to require that the loops enclosing an array are nested according to the hierarchical order of the array dimensions. Thus, the dimension being declared on the top level of the data structure will be scheduled on the outmost level. Because the range of a dimension is not allowed to depend on the subscript value of any lower level dimension in the data structure, in the example above when the loop of <A,3> is to be scheduled, the loops of <A,1> and <A,2> would have been scheduled on the outer levels. However, this requirement is unnecessarily

strong. For example, if we follow this scheme, then all the two dimensional arrays will have to be computed row-wise. With this restriction we may lose the opportunity to generate an optimal program.

A generalized solution would be to treat the range arrays as functions and find the real arguments of the range functions. For example, an N dimensional range array SIZE.X(I1,...,In) may be considered as a function which maps an N tuple of integers I1, ..., In to an integer value which is the range of the n+1th dimension of array X. Every subscript of the range array may be viewed as corresponding to an argument of the function. We will use the terms range array and _range_ _function_ interchangeably. Some of the function arguments may not affect the function value, namely the range does not vary with the value of these subscripts. The rest of the arguments which do play roles in determining the actual value are called _real_ _arguments_ of the range function.

By analyzing the assertion which defines a range array, we can find all the real arguments of the range array. If the range of a node subscript $<n,d>$ is specified by a range array and the range array has some real arguments, the real arguments of the range array should correspond to some other node subscripts of node n. In the generated program the loops which correspond to the real arguments should be scheduled on the outside level of the loop which corresponds

to the node subscript $\langle n,d \rangle$. For example, consider the specification in Fig. 5.2(a). The range array SIZE.A has two real arguments, i.e. $\langle SIZE.A,1 \rangle$ and $\langle SIZE.A,2 \rangle$. Since the node subscript $\langle A,3 \rangle$ references the range array SIZE.A and the node subscripts $\langle A,1 \rangle$ and $\langle A,2 \rangle$ correspond to $\langle SIZE.A,1 \rangle$ and $\langle SIZE.A,2 \rangle$ respectively, node subscripts $\langle A,1 \rangle$ and $\langle A,2 \rangle$ will be stored in the real argument list of node subscript $\langle A,3 \rangle$. It is shown in Fig. 5.3. The loop iterated on $\langle A,1 \rangle$ and $\langle A,2 \rangle$ will be scheduled on the outside of the loop on $\langle A,3 \rangle$. Similarly, we can find the real argument lists for $\langle al,K \rangle$ and $\langle B,3 \rangle$.
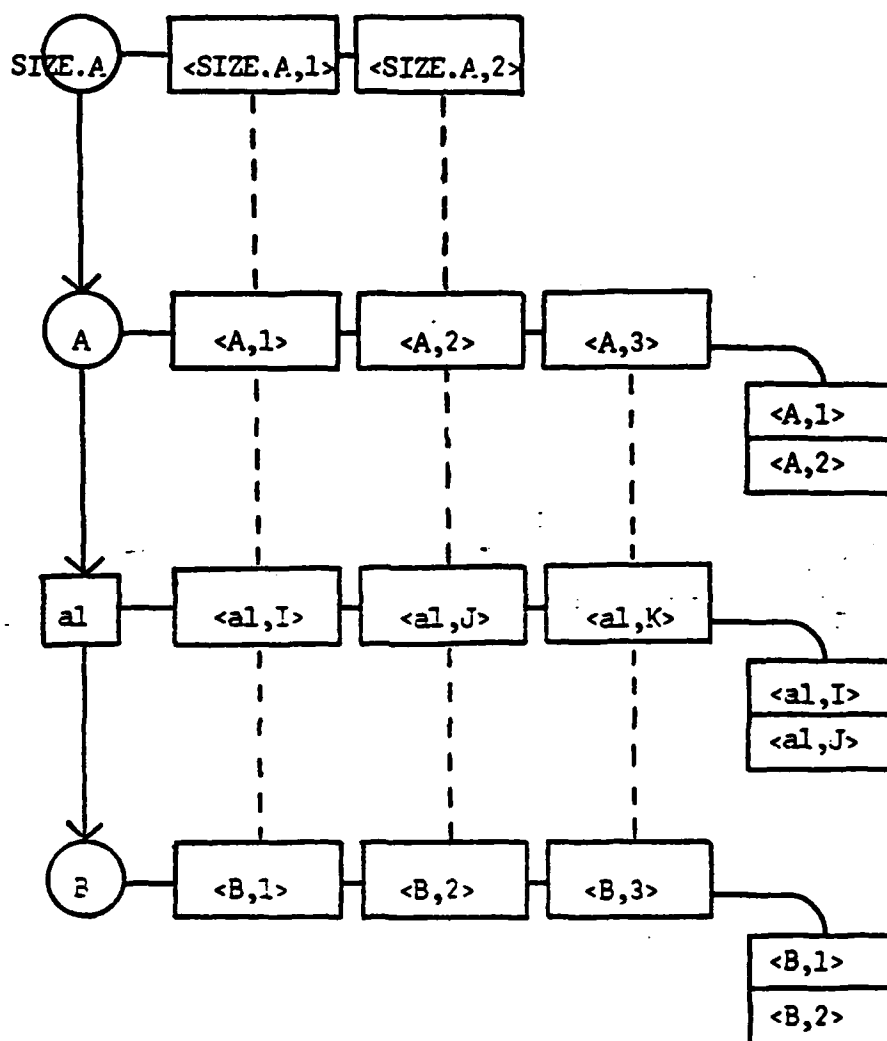
Fig. 5.3 Real argument lists of node subscripts

Example We will show how transposing an array effects the
mapping between the real arguments of the range arrays.
Let us examine the following assertions.

$$B(I,J,K) = A(J,I,K) \; ;$$

$$SIZE.A(M,N) = h(M,N) \; ;$$

Assuming that R(<A,1>) is equal to R(<B,2>) and R(<A,2>)

is equal to R(<B,1>). The range for subscript <B,3> is obtained from R(<A,3>) which is given by SIZE.A. SIZE.B(N,M) should be equal to SIZE.A(M,N). All we need is a permutation of subscripts to make the range array SIZE.A the same as SIZE.B. A possible flowchart for the loops enclosing node A and B is shown in Fig. 5.4.

```
        .
        .
DO <A,1> ;
  DO <A,2> ;
    DO <A,3>= 1 TO SIZE.A(<A,1>,<A,2>) ;
      A(<A,1>,<A,2>,<A,3>) ;
    END ;
  END ;
END ;
        .
        .
DO <B,1> ;
  DO <B,2> ;
    DO <B,3>= 1 TO SIZE.A(<B,2>,<B,1>) ;
      B(<B,1>,<B,2>,<B,3>) ;
    END;
  END ;
END ;
        .
        .
```

Fig. 5.4 Transposition of real arguments of
a range array

It should be noted that the order of the node subscripts <B,1> and <B,2> in the range array reference SIZE.A(<B,2>,<B,1>) is significant in the loop control statement for <B,3>. Therefore, in the real argument list associated with the node subscript <B,3> we should store the real arguments in the order of <B,2> followed by <B,1>. (Refer to Fig. 5.5)
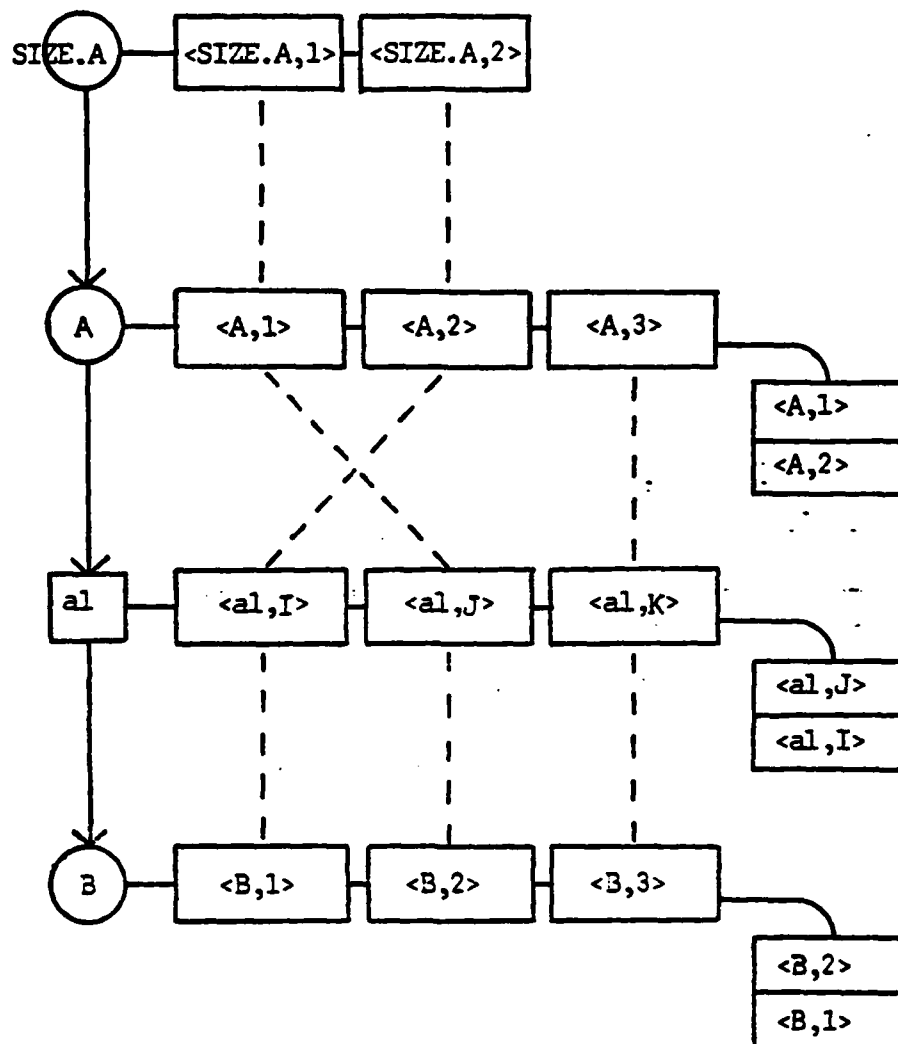
Fig. 5.5 The order of real arguments in the
real argument list

5.5 RANGE PROPAGATION ALGORITHM (RNGPROP)

The range propagation algorithm consists of three
steps. First of all, we locate the node subscripts which

have user specified ranges (Algorithm 5.1). In the second
step we propagate the explicit range specifications by
partitioning the node subscript set into range
sets (Algorithm 5.2). In the third step, we will propagate
the real argument list(RAL) among the node subscripts in the
same range set (Algorithm 5.3).

The data structure used are as follows. The total
number of node subscripts is denoted by $ALLSUBS. Every
node subscript is assigned a unique sequence number. A
vector TERMC(DICTIND) of integer denotes the kind of range
specification used for the least significant dimension of
each node. It can have the values of 1-4 to denote the
following conditions:

    1: the data structure has a constant number of repetition.

    2: the range is specified by an END array.

    3: the range is specified by a SIZE array.

    4: the range is implied by reading an end of file.

The vector LTERMC provides the same information for node
subscripts as TERMC for the nodes. The contents of TERMC
and LTERMC are computed by Algorithm 5.1.

<u>Algorithm 5.1 Find User Specified Ranges</u>

Output:

TERMC: The type of user specified range of every node in
      the Array Graph.

LTERMC: The type of user specified range of every node
       subscript.

1. Initialize the vectors TERMC and LTERMC to 0.

2. For each node n, in turn do:

    If attribute VARYREP=0, then TERMC=1.

    If attribute ENDB>0, then TERMC=2.

    If attribute SIZEB>0, then TERMC=3.

3. For every node n, in turn do:

    If TERMC(n) is not equal zero, find the node subscript <n,d> which corresponds to the least significant dimension of node n. Set the LTERMC entry of the node subscript to TERMC(n).

    Three arrays, HEADER, SETNEXT, and LRANGEP are used in step 2. Each of them has $ALLSUBS number of entries. HEADER(I) gives the sequence number of the header element of the block to which the Ith node subscript belongs. SETNEXT(I) links the Ith node subscript to the next node subscript in the same block, if any. When the Ith node subscript is the header of a block, then LRANGEP(I) shows the range of the Ith subscript. Algorithm 5.2 partitions the set of all the node subscripts. Initially every node subscript forms a block by itself. Then whenever we find that two node subscripts could have the same range and no range conflict would occur, we will merge their blocks. This merging process will continue until no further merging can be done. Since every node subscript can only be in one block at any moment, this is in fact a disjoint-set union problem[AHU 74]. The blocks formed in Algorithm 5.2 are

called range sets.

Algorithm 5.2 Propagation of Range Specification

Input:

LTERMC: The type of user specified range for every node subscript.

Output:

RANGE: A field in the LOCAL_SUB data structure of every node subscript. It contains the range set number where the node subscript belongs.

$RNGSET: The total number of range sets.

SET$RNG: The node number of the header of a range set.

Data structures:

$ALLSUBS: The total number of node subscripts.

HEADER($ALLSUBS): The node number of the header of the range set of a node subscript.

SETNEXT($ALLSUBS): For every node subscript, it points to the next node subscript of the same range set.

LRANGEP($ALLSUBS): If a node subscript is not the header of any range set, the value is -1. Else, if the node subscript has a user specified range, the value is the data node number of the range. Otherwise, the value is 0.

1. Initialization.

Make every node subscript a block by itself. For all values of I from 1 to $ALLSUBS do:

HEADER(I)=I,

SETNEXT(I)=0, /* NO NEXT ELEMENT */

LRANGEP(I)=node of the range /* IF IT HAS A DEFINED
RANGE */

=0, /* OTHERWISE   */

2. Merge blocks of the same global subscript name:

For every node subscript with sequence number I, check
whether it has a global subscript name. If it is a
global subscript of the form FOR_EACH.X or user declared
subscript X, let J be the sequence number of the node
subscript which is associated with the least significant
dimension of node X. Call procedure UNION(I,J) to merge
the blocks containing these two subscripts.

3. Propagate ranges between data nodes and control arrays
   or target nodes and assertion nodes:

For every edge in the Array Graph with edge type not
equal to 3 check the type of the subscript expressions
associated with the edge.   These edges connect data
arrays to the associated control arrays and the assertion
nodes to their target variables. For every subscript of
the source node, find the corresponding subscript in the
target node. If the APR_MODE of the subscript expression
is 1 or 2, merge them using procedure UNION.

4. Propagate ranges from assertion to source variable:

Scan all the edges of type 3 which connect a source
variable to an assertion. The range is to be propagated
backwardly. If the subscript of the source node has a
defined range, no merge will be done. Otherwise check if
the APR_MODE of the subscript expression is 1 or 2.   If

yes, call procedure UNION to merge it with the corresponding subscript of the target node.

5. The same as step 4. Except that no merge will be done if the subscript of the target node has a defined range.

6. Check the header of each block. If it does not have a user defined range, check the elements of the block. If there exists an element which is associated with a data node at or above record level and being the rightmost node in an input file structure, we may use end-of-file as the default range.

7. Assign a _range_ _set_ _number_ to every block of the partition. If a node subscript belongs to the kth block, put k into the RANGE field in the data structure LOCAL_SUB of the node subscript. Also store the node number which gives the range information of the block in SET$RNG(k) entry.

Procedure UNION(I,J)

Input:

I,J: The subscript sequence numbers of two node subscripts for which the range sets will be merged.

Output:

Modify the data structure HEADER, SETNEXT, and LRANGE to reflect the merging of the two range sets.

1. If both subscripts I and J are in the same block, exit.

2. If the blocks containing subscript I and J have different ranges, exit.

3. Put HEADER(I) into A.

4. Put HEADER(J) into B.

5. Change the HEADER entries of all the elements in the same block as J to A.

6. Append the list with the header B to the list with the header A.

7. Replace LRANGEP(A) by LRANGEP(B) if LRANGEP(A)=0.

8. Set LRANGEP(B) to -1.

Step three examines all the range sets. If the range of a range set is specified by a range array, a RAL is computed for every node subscript in the range set.

Algorithm 5.3. Propagation of Real Argument List

Input:

LTERMC: Type of user specified range of every node subscript.

RANGE: A field in the LOCAL_SUB data structure of every node subscript. It contains the range set number where the node subscript belongs.

Output:

RALP: A field in the data structure LOCAL_SUB of every node subscript. For every node subscript whose range is of types 2, 3, or 4, it points to a list of real arguments of the range function.

Data structure:

The real argument list pointed to by RALP consists of a list of elements which are stored in the data structure RAL. The fields in the RAL are as

follows.

$RAL: The number of real arguments.

RSPOS($RAL): The subscript position of a real argument in the range array.

MSPOS($RAL): The subscript position of the corresponding real argument in the node subscript list.

1. For each node subscript which has a user specified range and the termination criterion is not constant, form the RAL for it and put it into a candidate queue. (Refer to Algorithm 5.4)

2. Iterate step 3 to step 7 until the candidate queue becomes empty.

3. Get a node subscript from the queue. Let it be the subscript S of node X. Propagate the RAL of S to other node subscripts in step 4, 5, 6, and 7. If any node subscript gets its RAL newly defined, put it into the candidate queue such that its RAL can be propagated to other subscripts.

4. For each outgoing edge from node X, propagate the RAL of subscript S from node X to the target node. (Refer to Algorithm 5.5)

5. For each incoming edge into node X, propagate the RAL of subscript S from node X back to the source node. (Refer to Algorithm 5.6)

6. If subscript S references a global subscript, propagate its RAL to the global subscript.

7. If subscript S is a global subscript, then propagate its

RAL to all the subscripts which reference its name.

8. Stop.

Algorithm 5.4. Find RAL from a range specifying assertion

Suppose the range of the subscript $\langle X,n \rangle$ is specified by an assertion. Let the range array be SIZE.X or END.X. The algorithm tries to find the RAL for subscript $\langle X,n \rangle$.

1. Put all the subscripts of the target variable of the assertion which defines the control variable SIZE.X or END.X into a list.

2. If the target variable is END.X, delete the subscript on its least significant dimension from the list.

3. Repeat for each of the subscripts in the RAL to check whether it is referenced on the right hand side. If yes, it is a Real Argument. Otherwise, delete it from the list.

4. The resulted list is the RAL of the subscript $\langle X,n \rangle$.

Algorithm 5.5. Propagation of RAL forward along an edge

Assume S1 is a subscript of node X and there is an edge E from node X to node Y. The algorithm propagates the RAL of S1 to some subscript of node Y.

1. If the subscript expression of S1 is not type 1 or type 2, exit.

2. Let the corresponding subscript of node Y be S2. If RAL of S2 is defined, exit.

3. If the ranges of S1 and S2 are different, exit.

4. For each subscript in the RAL of S1, check its subscript expression type. If any one of them is not type 1, exit. Find their corresponding subscripts in node Y and form a new list. If the ranges of the corresponding subscripts are not the same, exit.

5. The newly formed subscript list is the RAL of S2.

Algorithm 5.6. Propagation of RAL backward along an edge

Assume S1 is a subscript of node X and there is an edge E from node Y to node X. The algorithm propagates the RAL of S1 to some subscript of node Y.

1. If there is no subscript of node Y corresponding to subscript S1, exit.

2. Let the corresponding subscript of node Y be S2. If RAL of S2 is defined, exit.

3. If the ranges of S1 and S2 are different, exit.

4. For every subscript Xi in the RAL of S1 find its corresponding subscript Yj of node Y.

    4.1 Let the subscript position of Xi in the local subscript list of node X be i.

    4.2 Check the LOCAL_SUB$ field in the data structure EDGE_SUBL associated with edge E. If the jth LOCAL_SUB$ is equal to i, the jth node subscript Yj in the local subscript list of node Y corresponds to Xi.

    4.3 Check the APR_MODE corresponding to subscript Yj in edge E. If it is not 1, exit.

4.4 Check the RANGE field of the node subscript Yj and that of subscript Xi. If they are different, exit.

5. Form a subscript list which contains those subscripts Yj's of node Y. It is the RAL of subscript S2.

Algorithm 5.7. Propagate RAL between Global subscripts

Suppose subscript S1 of node X and subscript S2 of node Y have the same global subscript name. The algorithm propagates the RAL of S1 to S2.

1. If the RAL of S2 is defined, exit.

2. For each subscript T in the RAL of S1, get its range, say RT. Check all the subscripts of node Y. If there is one and only one subscript U which has the same range as subscript T, then subscript U is the corresponding subscript of T. Otherwise, exit.

3. Form a subscript list which contains those subscripts U's of node Y. It is the RAL of S2.

5.6 DATA DEPENDENCY OF RANGE INFORMATION

In section 4.4.2 we have mentioned that range arrays cause implicit data dependency relationship. The edges of type 13 and 14 in the Array Graph represent this type of data dependency. However, it is not enough if we only have the edges from a range array SIZE.X or END.X to the node X. For every node in the Array Graph, no matter whether it is a data or an assertion node, as long as one of its node

subscripts is in a range set where the range is deined by a range array, an edge should be drawn from the range array to that node.

We can tell the range of every node subscript only after the range propagation phase. Therefore, the correct time to add this type of data dependency relationship is after we have found all the range sets. If a range set has a range array as its range specification, then there will be edges emanating from the range array and terminating at every node in the range set. Subscript expressions of type 1 are associated with the edges emanating from a SIZE range array. Subscript expression of type 2 is associated with the least significant dimension of an END range array and type 1 subscript expressions are associated with the other dimensions of the END range array.

CHAPTER 6

SCHEDULING

6.1 OVERVIEW OF SCHEDULING

Through the phases of data dependency analysis, dimension propagation, and range propagation we have analyzed the user's specification and checked the consistency and completeness of the specification. In a non-procedural programming language, the execution sequence is not specified in the program specification. The objective in this chapter is to determine the order of execution in performing the specified computation. We have collected the needed information in the convenient form of the Array Graph. The Array Graph contains all the program activities as nodes and the data dependency relationships as edges. The next step toward constructing a program is ordering the program activities represented by the nodes of the Array Graph under the constraints posed by: a) the edges of the Array Graph, and b) considerations of computation efficiency. As stated in chapter 1, efficient

scheduling is one of the main contributions of the reported research. This method of synthesizing the program is called scheduling here. It is followed by the actual program code generation.

Two rules which are frequently accepted in programming, except in cases where memory limitations are extremely severe, will be followed here as well. The first is that every input file is to be read only once. This rule will reduce the number of input activities which are usually relatively slow. If necessary we may store the input data in the memory for repetitive use. However, sometimes the memory price may be very high due to the large capacity of external storage. The second rule is that no values are to be recomputed. This means that once an element has been computed it will be retained as long as it is needed for later reference.

6.1.1 A BASIC APPROACH TO SCHEDULING

A correct but often inefficient realization of a computation can be obtained through the f·llowing scheduling method. Our eventual approach will be partly based on this simpler basic approach. The acyclic portions of an Array Graph may be scheduled very simply as follows. A topological sort algorithm can be applied to obtain a linear ordering of the nodes in the graph in accordance with the

edge constraints. Multi-dimensional nodes are then enclosed within nested loop controls. Every loop iterates the respective node over the instances of one of the distinctive node subscripts of the node.

When there are cycles in the Array Graph, a topological sort will not succeed. Superficially, a cycle in the Array Graph means a circular definition which does not allow us to determine a linear order for the computation. Actually since the Array Graph masks some of the details of the relationships in the corresponding Underlying Graph (see Chapter 4), there may be a cycle in the Array Graph where there are no cycles in the corresponding Underlying Graph. Also iterative solution methods can be applied to perform the computations even where there are cycles in the Underlying Graph. We have to apply a deeper analysis of the nodes and subscript expressions used in assertions in the cycle. The cycles that are found to be really not circular can be resolved to generate a linear schedule. The method employed is briefly described as follows. The Array Graph is decomposed into subgraphs. Each subgraph is a most strongly connected component (MSCC). A MSCC in a directed graph is a maximal subgraph in which there is a path from any node to any other node. The deeper analysis is then applied to the MSCC components in the Array Graph. The analysis described in section 6.2 consists of search of a dimension that is common to all the nodes in the MSCC. If

an edge is found in the MSCC which has an I-k type subscript
expression associated with it, the edge may be deleted.
This sometimes results in an acyclic subgraph which can be
topologically sorted. If this method is not successful then
other analysis methods, or alternatively an iterative
solution method may be applied.

6.1.2 EFFICIENT SCHEDULING

In general, a schedule which satisfies the constraint
of the data dependency relationship is not unique, if one
exists. Therefore, there is a degree of freedom to select a
schedule which meets efficiency requirements as well. We
want to have a schedule with the fewest number of loops or
with the least amount of working storage for the program
variables. Although we will use here the results of the
basic scheduling approach mentioned above, our method of
scheduling consists essentially of a process of repeated
merging of basic MSCCs in the Array Graph. As will be
shown, in this way we can reduce the use of memory and
computation time.

Non-procedural programming uses as many variables as
the values that occur during the program computation. If we
simply allocate separate memory space to each variable, as
may be done in the basic approach, we will most probably get
a program which uses a large amount of memory space and in

some cases may not be executable. Therefore, we are here primarily concerned with memory efficiency of the program. Our approach is to examine the effect on use of memory due to merging of blocks of nodes of the same or related subscript ranges and form iteration loops for the selected subscripts enclosing the merged blocks. We will select mergers of blocks of nodes which reduces the use of memory the most.

In some cases we have an alternative of maximizing the scope of one loop at the cost of reducing the scope of one or more other loops. The choice of which loop scopes are maximized is based on comparison of memory requirements of the alternatives. The alternative that requires least memory space for program variables will be selected.

The repetitions indicated by the node subscripts are controlled by loop statements. The execution of loop statements takes some CPU time. If the loop scopes in a program are small, i.e. if they contain fewer nodes, then there will be more loops in the program and the overhead spent on the loop control statements will be increased. This is another reason why it is desirable to maximize the loop scopes in the generated programs.

6.1.3 OUTLINE OF THE CHAPTER

The material in sections 6.2, 6.3, and 6.4 forms a background to understanding the optimization in the scheduling algorithm. In section 6.2 we will discuss the analysis of MSCCs. The algorithm of our optimizing scheduler is based on deeper analysis of cycles. A similar approach was used previously in an earlier version of the MODEL processor. Some changes discovered in the course of the presently reported research have been added. The merger of components is discussed in section 6.3. There are two bases for merging of components: when components have the same subscript ranges and when they have related range (this is explained later). In section 6.4 we will introduce the memory penalty concept which will be used to evaluate the use of memory in a partially designed schedule. The memory penalty is the memory cost associated with a candidate subschedule. The scheduling algorithm is presented in section 6.5.

6.2 ANALYSIS OF MSCC

6.2.1 CYCLES IN THE ARRAY GRAPH

A cycle in the Array Graph means that a variable definition depends directly or indirectly on itself. An Array Graph is a compact representation of an Underlying

Graph. It does not show the details of precedence
relationships in the Underlying Graph. Therefore, the
apparent circularity may be deceptive and not be reflected
in the Underlying Graph. In this case a correct computation
may be realized for an Array Graph cycle.

Consider for example the assertion in Fig. 6.1 which
defines the factorial function. Because of the recursive
definition there is a cycle in the Array Graph. But there
is no cycle of precedence relationship in the corresponding
Underlying Graph. Therefore, there exists a precedence
ordered sequence for computing all the factorial values.

a(I): F(I) = IF I=1 THEN 1 ELSE I*F(I-1) ;

(a) Assertion
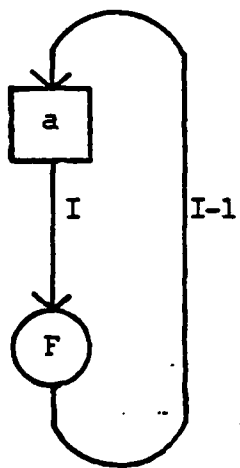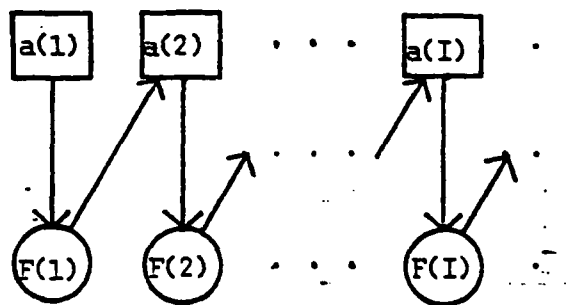
(b) Array Graph   (c) Underlying Graph

Fig. 6.1 Example of cycles in the Array Graph

A MSCC in the Array Graph may or may not represent a circular definition. If it is not truly circular, we may be able to perform the respective computation by using an iteration loop. In section 6.2.2 we will discuss the conditions under which a MSCC can be enclosed in a loop. If these conditions are met, we will find the loop parameter to bracket the entire MSCC. Once such loop is found, since the loop indices are ascending, the precedence relationships between the respective loop instances is assured. Therefore, as shown in section 6.2.3 we delete edges with I-k subscript expressions and the MSCC may be decomposed. If the above method fails, there are other approaches to schedule a MSCC which will be discussed in section 6.2.4.

6.2.2 ENCLOSING A MSCC WITHIN A LOOP

The objective of iterative computations of a single data or an assertion node is to define all the elements corresponding to the values of node subscripts associated with the node. In general, the values of every node subscript can be stepped independently of other node subscript values. Therefore, a node with N node subscripts would have an N level nested loops enclosing it, and each level of the nested loop corresponds to one distinctive node subscript. We will associate with every loop a loop variable with values which are stepped up by one from one to

the upper bound of a subscript range. All the nodes inside the scope of a loop will be executed once for every possible value of the loop variable. Generally if a node does not have a node subscript corresponding to a loop variable, the repetition would be redundant. We want to treat an entire MSCC in some manner as a single node, i.e. to compute all the elements of the nodes in the MSCC iteratively. We require however that all the nodes of a MSCC have a node subscript with which a loop brackets the MSCC. If one of the nodes does not have such a node subscript then the activity represented by the node, such as input/output, may be repeated, which will cause an erroneous computation. All the distinguished dimensions must then have the same range. It should be noted that the loop variable is stepped up each iteration by one, and no computation of a loop instance can depend on any computations in later loop instances.

Given a MSCC in the Array Graph, we will first check if all the nodes in the MSCC have more than zero dimensions. If every node does have at least one dimension to schedule, we will then check the subscript expressions on the edges of the MSCC to see if the entire MSCC can be enclosed within a loop. The edges in the Array Graph represent relationships between some elements of the nodes at the ends of the edges. The subscript expressions associated with edges reveal more precisely the precedence relationships between specific elements. In the following we examine the subscript

expressions associated with an edge to determine if the nodes at the end of the edge can be scheduled within the scope of a loop.

Definition Let A be a node of n dimensions. Then $\underline{A}$ denotes the set of all the instances of node A, i.e. $\underline{A} = \{A(I1,\ldots,In)\mid 1<=Ik<=R(<A,k>), \text{ for } 1<=k<=n \}$.

Definition Let A be a node of n dimensions. Then $\underline{A}(Ii=C1; Ij=C2; \ldots)$ denotes the set of all the instances of node A with the ith subscript Ii being C1 and the jth subscript Ij being C2, ... etc.

Consider an edge from node $A(J1,\ldots,Jm)$ to node $B(I1,\ldots,In)$ in the Array Graph:

$B(I1,\ldots,Ik,\ldots,In) \longleftarrow A(E1,\ldots,Ep,\ldots,Em)$

where J's and I's are the node subscripts of node A and B respectively, and E's are the subscripting expressions of A. Consider the subscript expressions of types 1, 2, 3, and 4.

1) If a subscript expression Ep is of type 1 and equals to Ik, then every element in $\underline{B}(Ik=c)$ depends only on the elements in $\underline{A}(Jp=c)$. Since $\underline{B}(Ik=c)$ does not depend on any element in $\underline{A}(Jp=d)$ with d>c, the Underlying Graph dependencies are satisfied if node A, followed by B, are bracketed by a loop where the parameters of the iteration are the pth dimension of A and the kth dimension of B. These are referred to as a distinguished dimension of A or of B.

2) If the subscript expression Ep is type 2 or 3 and equals to Ik-a, then for any positive integer c every element in $\underline{B}$(Ik=c) depends only on the elements in $\underline{A}$(Jp=c-a). Since the parameters of the bracketing loops are in ascending order (in step of 1) then this assures that $\underline{A}$(Jp=d) is computed before $\underline{B}$(Ik=c) with d<c. Thus it is allowed to schedule node A and B into one loop, with Ik and Jp the distinguished dimensions.

3) If the subscript expression Ep is type 4, then for any positive integers c and d every element in $\underline{B}$(Ik=c) may depend on elements in $\underline{A}$(Jp=d). We will be conservative and assume that every element in $\underline{B}$(Ik=c) depends on at least one element in $\underline{A}$(Jp=d) with d>c. Therefore, it is impossible to designate the pth dimension of A and the kth dimension of B as the distinguished dimensions for a loop.

Example Given an assertion al as follows. Let A and B be square arrays. There is an edge from array node A to assertion node al.

al(I,J): B(I,J) = A(g,J);

where g is a type 4 subscript.

Consider the node set {A,al}. Consider scheduling this set into one loop with <A,1> and <al,I> as their distinguished dimensions. Let SA be {A(J1,J2)|J1=2} and SB be {al(I,J)|I=1}. SB is in the first instance of the loop and SA is in the second instance of the

loop, therefore SB precedes SA. Consider next the element al(J of SB. We can find an element A(2,2) in SA which precedes al(1,2) because of the type 4 subscript on ⟨A,1⟩ dimension. SB and SA then precede each other, in the Underlying Graph, and therefore can not be scheduled.

Example Given the assertion a2 below.

$$a2(I,J): \quad Y(I,J) = X(I,J) + X(J,I);$$

X is a square array and subscripts ⟨X,1⟩, ⟨a2,I⟩, and ⟨a2,J⟩ have the same range. We want to schedule the node set {X,a2} in one loop with ⟨X,1⟩ and ⟨a2,I⟩ as the distinguished dimensions.

All the subscript expressions being used with node X are not type 4. However, in the term X(J,I) a subscript J occurs on the distinguished dimension of X, i.e. ⟨X,1⟩. Since ⟨a2,J⟩ does not correspond to the distinguished dimension of node a2, it may be scheduled in an inner level loop and iterates faster than ⟨a2,I⟩, therefore some array elements of X will be referenced before defined. Thus we should not form a loop with these designated distinguished dimensions.

From the examples above we know that the subscript expression on the distinguished dimension of a node must not be a general expression and it should correspond to the

distinguished dimension of another node in the same loop, otherwise the loop can not be formed. Since the loop instances are strictly running upward starting from one and all the subscript expressions on the distinguished dimensions are of the form I or I-k, no reference goes to the later loop instances, therefore, no data dependency relationship is violated. In fact, by constructing the loop we have divided the whole computation into many smaller tasks where every task corresponds to a loop instance. It should be noticed that the formation of an outer loop does not exclude the possibility that the original computation involves an unsolvable cycle. What we are assured is that the outer loop divides the original problem into smaller ones and which can be solved easier.

6.2.3 DECOMPOSING A MSCC THROUGH DELETION OF EDGES

Consider now the case where an MSCC is scheduled in one loop based on the tests described in the previous subsection. The nodes in the MSCC have each a distinguished dimension which corresponds to the loop variable. Also the subscript expressions associated with the distinguished dimensions are of the form either I or I-k. We will show in the following that where the parameter of the loop is stepped up from one by a step of one then edges which have a subscript expression of type 2, i.e. I-k, are superfluous

and can be removed.

Consider an edge of the form $B(\ldots,I,\ldots)$ $\longleftarrow$ $A(\ldots,I-k,\ldots)$ where $I-k$ and I occur on the pth and the qth dimension of nodes A and B, respectively. If node A and B are scheduled in the loop of I, then the elements in $\underline{A}(Jp=I-k)$ have been evaluated in the I-kth loop instance and the elements in $\underline{B}(Iq=I)$ are evaluated in the Ith loop instance. Since the values of loop variables are ascending, therefore every element of $\underline{A}(Jp=I-k)$ precedes all the elements of $\underline{B}(Iq=I)$. This implies that the precedence relation represented by the above edge is superflous as it is enforced by the order of evaluation of the respective elements. In short, when two nodes are scheduled in a loop of loop variable I, the precedence relationship presented by subscript expression $I-k$ is subsumed by the order of loop execution. This is illustrated in Fig. 6.2, showing the Array Graph of a Factorial function which is defined with recursion. The recursion causes a cycle of two nodes {al, FAC}.

al: FAC(I) = IF I=1 THEN 1 ELSE I*FAC(I-1) ;
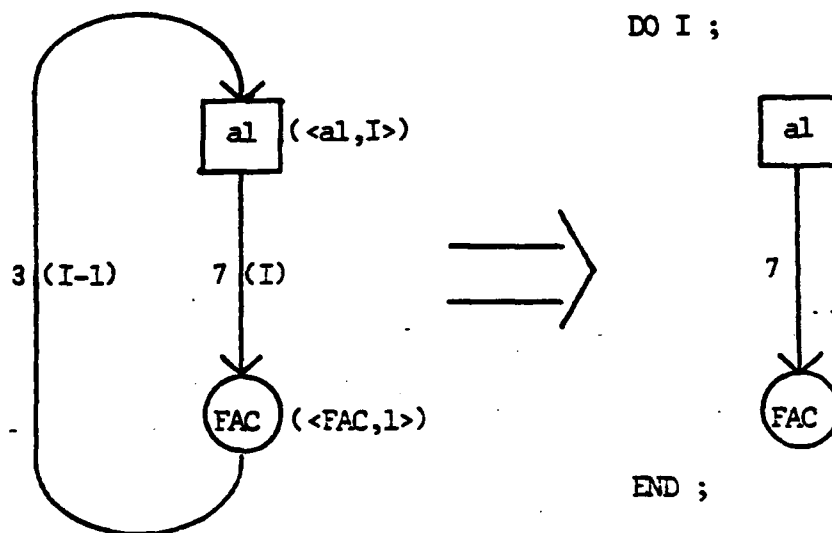
Fig. 6.2 Remove I-k edges in a loop

These two nodes can be scheduled in a loop iterating over node subscript <al,I>. The kth instance of the assertion al is evaluated in the kth loop instance and it references the k-1th instance of the array FACT, which has been evaluated previously in the k-1th loop instance. Therefore the edge associated with subscript expression I-1 can be removed. There is no further a cycle in the Array Graph.

6.2.4 OTHER APPROACHES TO DECOMPOSING AN MSCC

There are a number of methods for scheduling a MSCC in an Array Graph. We have been primarily interested in the cases that a cycle can be implemented by a loop with the parameter that runs upward from one. However, not all the cycles can be implemented with this simple loop mechanism. Thus if the above approach fails it will be necessary to apply other methods. Consider first the case where the array elements may be evaluated in a sequence which does not follow the natural ascending order of subscripts. Consider for example the following specification which defines A, a vector of 50 elements.

Example

```
A(I) = IF I=25 THEN X

        ELSE IF I<25 THEN A(I+2)+X

        ELSE A(I-1)+A(I-25) ;
```

A possible PL/I program to compute array A is as follows.

```
A(25) = X ;
DO I = 23 TO 1 BY -2 ;
  A(I) = A(I+2)+X ;
END ;
A(26) = A(25)+A(1) ;
DO I = 24 TO 2 BY -2 ;
  A(I) = A(I+2)+X ;
```

```
END ;

DO I = 27 TO 50 ;

  A(I) = A(I-1)+A(I-25) ;

END ;
```

When the subscript expressions are first order polynomials, we can divide an array nodes into many parts and compute the parts of the array separately [SHAS 78].

A cycle in the Array Graph may also be considered as a set of simultaneous equations and numerical methods such as Jacobi and Gauss-Seidel iterations can be applied to solve the system of equations [GREB 81]. Since splitting nodes in the Array Graph, as suggested by Shastry, is complicated to apply, the MSCCs which can not be decomposed may be treated similar to simultaneous equations and solved iteratively. In this dissertation we will refer only to the cases that a MSCC can be decomposed as described above. The other methods are described in the references.

6.2.5 A SIMPLE SCHEDULING ALGORITHM

The methods of scheduling an MSCC in a loop and attempting to decompose a MSCC may have to be applied repeatedly, depending on the outcome of each application. This section describes a simple scheduling algorithm which incorporates repeated application of the methods described

earlier. It generates a correct schedule based on an Array Graph. However it does not include the consideration of program efficiency.

The algorithm consists of two mutually recursive procedures, SCHEDULE_GRAPH and SCHEDULE_COMPONENT. Given any Array Graph as input, SCHEDULE_GRAPH procedure finds the MSCCs in the Array Graph. The MSCCs are then sorted into a sequence {M1,M2,...,Mn} which retains the partial order of the precedence relationships between the MSCCs. SCHDULE_COMPONENT procedure then schedules each component separately. If Si is the schedule of component Mi, the sequence {S1,S2,...Sn} is returned as the schedule of the original graph.

The input to procedure SCHEDULE_COMPONENT is an MSCC, say Mi. If Mi is a single node component and there is no unscheduled node subscript associated with it, the node itself is returned as the schedule of the component. Otherwise, the component may be schedulable in a loop. The procedure tries to find a loop variable which satisfies the requirements discussed in the previous section. If a loop variable is found, say I, it then deletes the edges in component Mi with subscript expression I-k and marks the distinguished dimensions of the nodes in Mi as scheduled. Let Mi' denote the resulting graph. Then it calls the procedure SCHEDULE_GRAPH to produce a schedule for the graph Mi'. After SCHEDULE_GRAPH returns the schedule of Mi', a

loop with loop variable I and loop body, the schedule of Mi'
is formed by SCHEDULE_COMPONENT and returned as the schedule
of Mi. If no loop variable can be found, SCHEDULE_COMPONENT
sends a warning message to the user and calls the procedures
described in section 6.2.4 to decompose the MSCC.

6.3 MERGER OF COMPONENTS TO ATTAIN HIGHER EFFICIENCY

The basic scheduling algorithm, described above,
consists essentially of topological sorting of the nodes or
MSCCs in the Array Graph and of the enclosing of these
entities within the scope of nested loops for the respective
dimensions. In contrast, the scheduling algorithm offered
here considers the Array Graph globally and progressively
merges components into the scope of a selected loop which
reduces the most the use of memory and computing time. The
scope of the loops in the schedule is thus progressively
enlarged.

Given an Array Graph as input, we can construct a
component graph where every MSCC is a component node and an
edge is drawn from component A to component B if and only if
there exists an edge in the original Array Graph which leads
from a node in the component A to a node in the component B.
The component graph is an acyclic graph. Note that the
MSCCs in an Array Graph are not further divisible. The
merger process starts with the MSCCs in the Array Graph as

the basic components, and through merger it creates larger components progressively. A loop scope can be the union of some MSCCs. In this section we will discuss the merging of MSCCs in an Array Graph into the scope of one loop.

6.3.1 MERGER OF COMPONENTS WITH THE SAME RANGE

The condition for scheduling a set of component in one loop is that every component in the scope of a loop have a distinguished dimension corresponding to the loop variable. There are several condition on designating distinguished dimension of a node in an Array Graph or a Component Graph. First the distinguished dimensions of the components must be in the same range set and have a common range which specifies the number of iterations of the loop. The loop variable is stepped up by one in successive iterations. Therefore also the order of execution of elements of each component will be evaluated in this order. The second condition is that an evaluation of each instance of a component in a loop instance should not refer to values computed in later loop instances.

Further, components to be merged into the scope of a loop may not depend on any other component which does not have a distinguished dimension and which in turn depends on one of the components to be merged. The rule is that a set of components which can be scheduled in one loop should be

equal to its <u>closure</u>. The closure of a set of components includes all the components which are reachable from any component in the set and which also reach any component in the set. For example, consider the component graph in Fig. 6.3. The components C1, C2, and C4 have a common dimension I. Still they can not be merged into the scope of a loop with the loop variable I. The closure of the set of components {C1, C2, C4} includes component C3. Since C3 does not iterate with subscript I, it can not be scheduled in the loop of I. Component C4 can be scheduled only after component C3. Therefore, at most we can merge components C1 and C2 or C2 and C4 into the scope of a loop.

The set

C1 I

The closure
of the set

C2 I

C3

C4 I
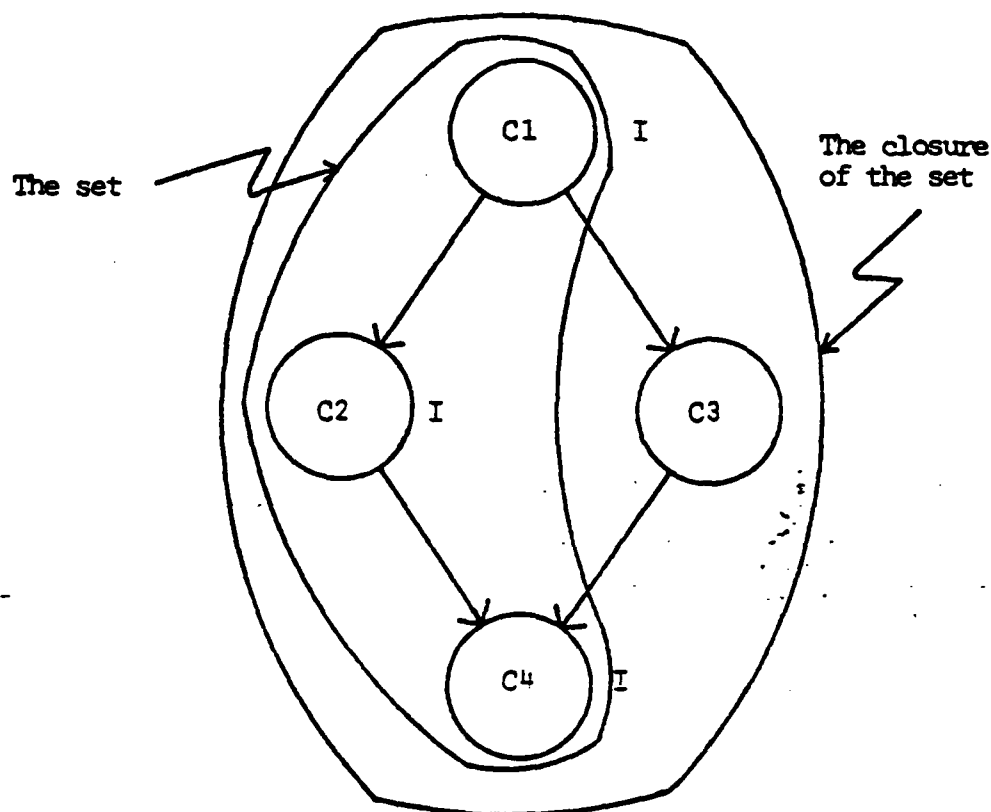
Fig. 6.3 Closure of a set of components

The search and selection of a distinguished dimension
for each component in a set is similar to the analysis of
subscript expressions in MSCCs described in section 6.2. We
showed there that the subscript expressions associated with
edges terminating at a component can not be type 4 and that
subscript expressions associated with the edge should
connect the distinguished dimensions of the components at
the ends of the edge.

6.3.2 MERGER OF COMPONENTS WITH SUBLINEARLY RELATED RANGE

In the previous subsection, we considered merging components with distinguished dimensions which have exactly the same range as the loop variable. Every node is then executed once in each loop instance.

There is a large class of cases where subscript expressions are explicitly related, i.e. where we use an indirect subscript $X(I)$ and X is a function of I. Statements with such an indirect subscript may in some case be conditionally executed in the scope of a loop for the parameter I. We will require that the indirect subscript expression $X(I)$ have values which grow monotonically and slower than that of the loop variable I. This feature of sublinearity was already mentioned in section 4.4.2. As explained in [PNPR 80], use of indirect sublinear subscript is important in many instances, such as selecting a subset of records from a sequential file or merging two sequential files into one.

In section 4.4.2 we have discussed the criterion for recognizing a vector which can be used for indirect indexing. The values of elements of an indirect indexing vector grow slower than the subscript value of the elements. The range of its dimension will be called here the major range, while the range of its content will be called subrange relative to the major range. For example, the

variable X in Fig. 6.4 satisfies these criteria. X is used in the subscript expression of the first dimension of node A and therefore $R(\langle X,1 \rangle)$ is a major range and $R(\langle A,1 \rangle)$ is a subrange relative to $R(\langle X,1 \rangle)$.

```
X(I) = If I=1 THEN 1
            ELSE IF <condition is true> THEN X(I-1)+1
                                        ELSE X(I-1) ;

B(I) = A(X(I)) ;
```

Fig. 6.4 Example of indirect sublinear indexing
in subscript expression

A subrange relative to a major range may be the major range of some other subranges. Therefore, the sublinear relationship between the ranges may form a tree with the maximal major range at the root. We merge major ranges and subranges in a bottom up order. By progressively merging each subrange with the next level major range finally we will obtain a loop which iterates in the maximal major range, and where all of its subranges are nested inside the loop. Such merger of subranges may not always be possible. For example, if type 4 subscript expression is used in the distinguished dimensions of a component, the precedence relationship will prevent us from scheduling this component into the scope of a loop.

When a set of components with a subrange and a major range are merged into the scope of a loop, the major range will be used as the loop range and the value of elements of the indirect indexing vector will be checked to evaluate only the elements which are within the subrange. An instance of the subrange is executed for each stepping up by 1 of the indirect indexing vector. The computation of the indirect index should precede the computation of any node within the subrange. This introduces an additional precedence relationship.

We will treat subscript expressions of types 5, 6, and 7 similar to types 1, 2, and 3, respectively, in checking the consistency of subscript expressions of the distinguished dimensions as discussed in section 6.2.1. If a check of the subscript expressions of the distinguished dimensions fails, i.e. some type 4 subscript expressions are used or the subscript expressions do not connect distinguished dimensions of the components, we will treat these indirect subscript expressions of type 5, 6, and 7 as type 4. If the check succeeds, we will add edges in the Array Graph from the indirect indexing vector to the nodes referencing it. This is similar to the addition of edges from a range array to the nodes referencing the range array.

6.4 MEMORY EFFICIENCY

In some cases the same memory space may be shared by a number of variables, thereby using memory storage more efficiently. Small savings of memory space are not worth the cost of the analysis. For example, sharing memory space among few scalar variables does not save much memory space. Our approach will concentrate on having elements of the same array share the memory space. Since the range of each array dimension is in general large and there are several dimensions, the saving should be considerable. It should also be noted that memory space is statically allocated to the variables in the produced program. Compared with dynamic memory allocation, static memory allocation has the advantages of simplifying the program control in that there is no need to allocate memory space at run time. This also facilitates efficient random access of array elements.

Three alternative approaches to allocating memory are used:

1. **Physical Dimension**

   If all the elements along some array dimension have different memory spaces assigned to them, the memory space allocated is proportional to the range of the array dimension. This method of allocating memory will be referred to in the following as the **physical dimension**.

2. **Virtural Dimension**

If all the elements along some array dimension share the same memory space, a single element memory space serves for the entire array dimension. We will refer to this method of allocation as <u>virtual dimension</u>.

3. <u>Window of width k</u>

In some cases there is no need to store all the elements in an array dimension in main memory. But an array reference of the form A(I-k) makes it necessary to keep k+1 array elements in main memory at any moment. This type of memory allocation will be referred to as <u>window of width k+1</u>.

For every array dimension we have to decide how the memory space is to be allocated. The memory allocation decision is related to the program execution sequence. Different program schedules may require different memory allocation approaches. For example, Fig. 6.5 shows two different schedules for copying a file. The one which reads all the records into the main memory then writes them out takes more memory space than the other one which copies the file, record by record.
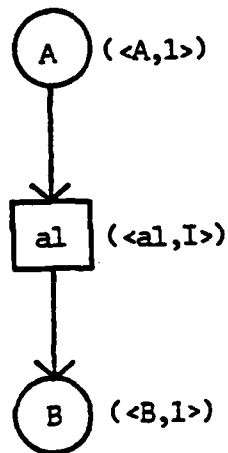
Fig. 6.5 Two schedules for copying a file

In the following we will show how the memory allocation decisions are influenced by the program schedule and how the memory space requirement for the program variables is evaluated.

6.4.1 EVALUATION OF MEMORY USAGE

We will first consider in what units we should allocate memory space. If a data structure or substructure is used as an argument of a function or an operation, the whole structure must be passed between program modules. The relative position of its constituent elements becomes important to the computation. Therefore we can not allocate memory space to its elements separately. On the other hand, economic allocation of memory space requires that the unit be as small as possible. We will require that all the operations operate on fields. Operations on higher level structure must be therefore transformed into operations on elementary data structure. The memory space will therefore be allocated in the unit of fields.

The array dimensions above the unit data structure will be considered as logical array dimensions for which there may not be corresponding physical dimensions in the allocated memory space. One of the three approaches mentioned above may be used to allocate memory space. Since a virtual dimension requires less memory space than a physical dimension, we would not physically allocate memory space to an array dimension unless it is necessary based on the logic of the specification. In the following we will discuss the conditions when an array dimension has to be physical or window of width k.

The values of data structures may be produced by some program activities such as reading an input file or evaluating an expression, and consumed by some other activities such as writing an output file or referencing an expression. If the production and consumption of the elements along an array dimension does not proceed in a planned order then all the array elements that are produced can not be discarded. All must be stored simultaneously in main memory.

Given a program schedule we can check whether the program activities which produce or consume the values along an array dimension are all in one loop. If not, that array dimension should be a physical dimension. If all the definitions and references of an array are in the same loop, we should further check whether any type 2 or 3 subscript expressions are used, because the occurrence of I-k type subscript implies the necessity of keeping previous k elements while computing a new array element. Thus the memory space for the array dimension should be a window of width k+1. It should be noted that if an array has its distinguished dimension using either a finite window or a physical dimension memory allocation scheme, all the loop for array dimensions which are scheduled nested inside the current loop have to be of physical dimensions. This is illustrated in Fig. 6.6, where a two dimensional array A is computed by a nested loop. Suppose the outer loop iterates

over the first dimension of A, i.e. <A,1>. The presence of subscript expression I-1 requires a memory allocation scheme of window of width two for <A,1> dimension. Since the array element of A is computed row by row and the computation of array elements in one row depends on the value of array elements in the previous row, therefore, we will have to allocate two rows of memory space for array A.

al: A(I,J) = IF I=1 THEN f(J)
                   ELSE g(A(I-1),J) ;

(a) MODEL specification

```
DO I ;
  DO J ;
    al(I,J) ;
  END ;
END ;
```

Array A

(b) Schedule
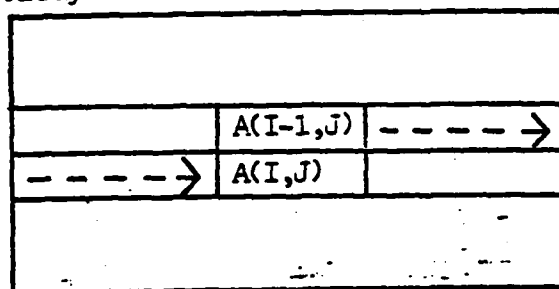
(c) Memory requirement

Fig. 6.6 Effect of window dimension on the outer loop
over dimensions on the inner loops

After the memory allocation approach for every array dimension has been determined, we can estimate the memory space requirement, which will serve as a measure of the program quality. Given an N dimensional array A, we can

define the required memory space M for a node subscript ⟨A,i⟩ as follows.

$$M(\langle A,i \rangle) = 1 \qquad \text{if the ith dimension is virtual,}$$
$$= k \qquad \text{if using window of width k,}$$
$$= \text{upper bound of } R(\langle A,i \rangle) \qquad \text{if physical.}$$

If an array dimension is not physical, the upper bound of its range is not used in calculating the memory requirement. The upper bound is needed to estimate the memory space for a physical dimension. Sometimes the range of an array dimension is specified by an assertion and the upper bound is not known until run time. In that case we can only assume the upper bound is infinity unless the user has specified an upper bound of the range in the data description statements. The memory space for array A is the product of $M(\langle A,i \rangle)$'s for all the dimensions of A. The total memory requirement of a program is the sum of the memory space used by every array variable.

6.4.2 MEMORY PENALTY

Analysis of the loop scope leads to the selection of the memory allocation scheme for the respective array dimension. The memory penalty of a loop is defined as the memory cost of the arrays included in the loop scope. The memory cost is the difference in memory requirements between the ideal case (virtual dimension) and the memory

requirements if the loop is formed. In order to evaluate the memory penalty of a loop, we first find all the nodes whose memory allocation scheme is influenced by the construction of the considered loop.

Whenever an Array Graph edge crosses the loop boundary, a source or target node of the nodes in the loop will be outside of the loop. Either one of the two nodes may require using the physical memory allocation scheme. For example, if an edge from a data node to an assertion node crosses the loop boundary, (i.e. the data node is in the scope of the loop while the assertion node is outside), the data node is defined in one loop and referenced outside it. Therefore, its array dimensions have to be physical. Similarly if the edge crossing the loop boundary is from an assertion node to a data node, the dimension of the target node has to be physical.

Each node under consideration may fall into one of the following three categories and the memory penalty can be computed accordingly.

1. A physical dimension for a distinguished dimension. This category is recognized by the existence of an edge which crosses a loop boundary. The memory requirement in ideal case is taken as that of a virtual dimension. The memory requirement for a loop is computed by multiplying the upper bounds of all the unscheduled dimensions and the dimension that is considered for a loop. The difference

is the penalty of the loop for this array.

2. A virtual dimension for the distinguished dimension. In this case the loop boundary is not crossed by edges and all the subscript expressions on its distinguished dimension are type 1 subscripts. The memory penalty for a virtual dimension should be zero.

3. A window of width k+1 for the distinguished dimension. Similar to the virtual dimension category. No edges would cross the loop boundary. However subscript expressions of the form I-k on its distinguished dimension are allowed. The other unscheduled dimensions are considered to be physical dimensions. The penalty is computed similar to the first category.

Example Consider the memory penalty of a loop shown in Fig. 6.7. The ranges of subscripts I and J are 10 and 20 respectively, and every data element occupies one unit of memory space. The memory requirements in ideal cases for node A, B, C, and D are 1, 1, 1, and 1 respectively. The memory requirements if the loop is formed will be 10, 40, 1, and 200 respectively. Arrays A and D have to be physical and the first dimension of array B needs a window of width 2. The memory penalty for this loop is the difference of 251 and 4, i.e. 247 units of memory space.
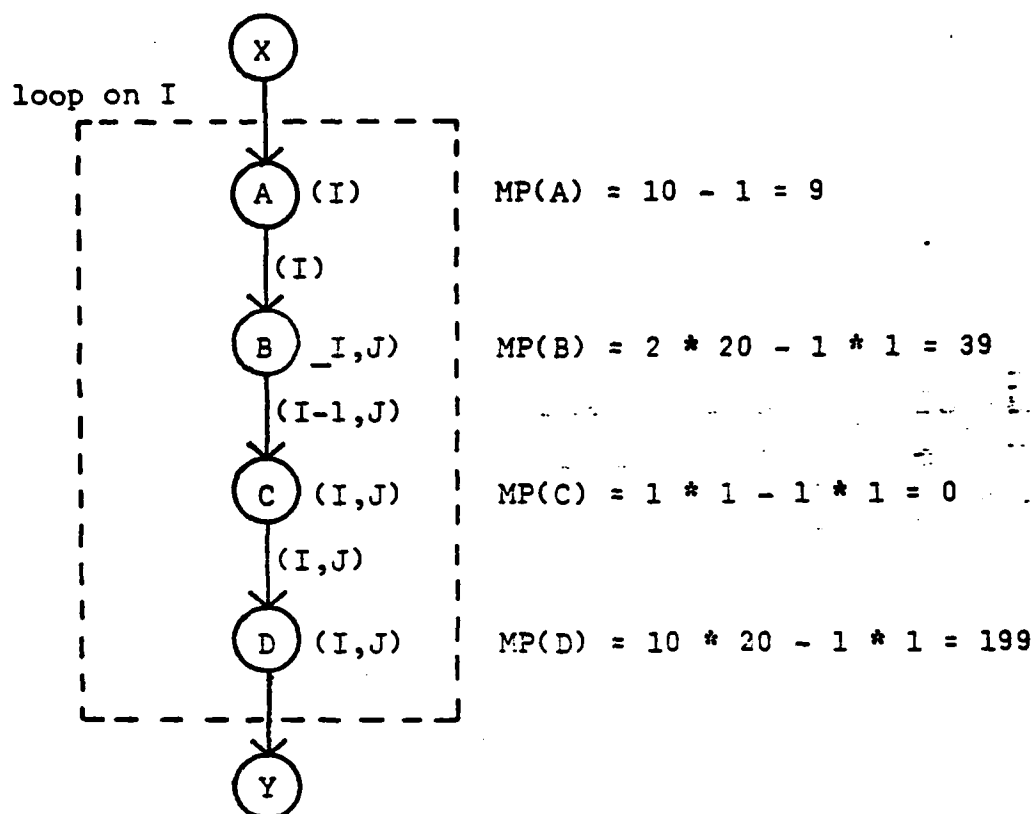
Fig. 6.7 Example of computing memory penalty

Information about the unscheduled dimensions may be used to compute the penalty more accurately. For example, some array dimensions must be physical dimensions because of the use of type 4 subscript expressions. During the process of scheduling, we can accumulate such information to speed up the memory penalty evaluations.

loop on I

X

A (I)   MP(A) = 10 - 1 = 9

(I)

B _I,J)   MP(B) = 2 * 20 - 1 * 1 = 39

(I-1,J)

C (I,J)   MP(C) = 1 * 1 - 1 * 1 = 0

(I,J)

D (I,J)   MP(D) = 10 * 20 - 1 * 1 = 199

Y

Fig. 6.7 Example of computing memory penalty

Information about the unscheduled dimensions may be used to compute the penalty more accurately. For example, some array dimensions must be physical dimensions because of the use of type 4 subscript expressions. During the process of scheduling, we can accumulate such information to speed up the memory penalty evaluations.

6.5 A HEURISTIC APPROACH TO MEMORY-EFFICIENT SCHEDULING

In general, there is a large number of schedules which can realize the computation of a program specification. The schedule with the minimal total memory requirement will be called an _absolute_ _optimal_ _program_. In principle it should be possible to enumerate all the possible schedules for an Array Graph, as there is a finite number of them, and then evaluate the memory requirement of each schedule. We would thus be able to find the absolute optimal schedule. For several reasons this method is not practical. The program events being scheduled are low level activities represented by nodes, i.e. statements and variables, and an Array Graph may easily consists of several hundred or even thousands of nodes. Also the nodes in the Array Graph may be multi-dimensional and the number of combinations of possible nested loops is very large. Further, the constraints on the feasible schedules are complicated. Thus enumerating all the feasible schedules would be prohibitive, and an exhaustive examination of all the feasible schedules to find the absolute optimum is not acceptable.

Instead we have adopted the heuristic approach as follows. Given an Array Graph as input, we first construct an acyclic component graph with the MSCCs in the Array Graph as nodes. Our objective is to repeatedly merge components in the component graph into blocks which correspond to loop scopes. This process will be applied repeatedly to the

levels of nested loops. On the first application it will produce the outer level loops. The blocks are formed by merging as many components as possible which have the same or related ranges. The process is repeated for each lower level of the nested loops, based on the subgraph that corresponds to the higher level loop. This process may not result in the absolute optimal program as the outer level loop scopes are determined without the analysis of the effects of inner loop structures on the use of memory space. However considering the effect of inner loops on memory usage is a complex process and it represents a large increase in the number of alternatives that must be evaluated. The scope of the major loops in a program are maximized in our proposed approach and there is no, or little, effect of inner loops on memory usage. Thus this heuristic approach represents a good compromise between the amount of analysis involved and the payoff in reducing memory usage.

On each level of loops, the scheduling process consists of a trial scheduling for every range set in the corresponding Component Graph. A loop for the range R will enclose only the components which have dimensions in the range set associated with range R. The range sets related to R (through sublinear indirect indexes) will later be merged with the blocks of range R. The maximum loop scope for every range R is the range set of R.

The trial scheduling of each range set consists of finding the closure of the range set and an attempt to schedule nodes in the set which may be within the scope of the respective loop. We first merge into a block the components in the range set which do not have any predecessors in the closure of the range set. Progressively we will merge into the block other components which depend on those in the block, as far as possible. The merger involves selection of a distinguished dimension in each component, as described above. At the end we evaluate the memory penalty of the loop scope obtained by the trial scheduling. The loop with the smallest penalty will be scheduled finally. This process will be repeated with the unscheduled portion of the graph until all the components in the Component Graph are scheduled.

There are many possible orders for merging components in the closure of a range set, to form the scope of a loop. For example, we may arbitrarily pick a component in the middle of the Component Graph and merge it with its neighbor components or start with a component on which no other components depend and merge the components backward. However, considering all the possible orders of mergers will further increase the number of alternatives that must be evaluated. The order of mergers is unimportant in the case where the whole range set can be scheduled in one loop, i.e. it is the case that all the array dimensions may become

virtual. No matter in what order we merge the components, we will finally get the same loop scope. Again, we selected the forward merging of the Component Graph as a good compromise between quality of the schedule and the amount of analysis.

It is necessary next to order the blocks associated with outside level loops in an execution sequence order. The memory cost will be the same for any order that maintains the precedence relations between these blocks. We choose to order the blocks by topological sorting. For every outer level loop we mark the distinguished dimensions of the blocks as scheduled.

We apply the scheduling algorithm recursively to each inner nested level loop by considering only the subgraph which contains the nodes in one loop scope. The resulting schedule will be the body of the outer level loop.

We will illustrate this process with an example of scheduling the Array Graph shown in Fig. 6.8. Every node is a MSCC by itself, and the initial Component Graph is in fact the Array Graph. The candidate ranges are $R(<A,1>)$ and $R(<B,1>)$. Assume that the repetition numbers are 500 and 200, respectively. The range set of $R(<A,1>)$ contains three nodes: A, a1, and C. The closure of {A, a1, C} is itself. If we schedule the whole set into one loop, the penalty will be making array B physical. On the other hand, the trial

scheduling of the range set of R(<B,1>) contains two nodes: B and a1. If this set is scehduled in one loop, the penalty will be making both array A and C physical. We will select the loop of R(<B,1>) since the size of array B is greater than the sum of the sizes of array A and C. We mark the component B and a1 as scheduled. There are two components left to be scheduled. We have no alternative but to schedule each of them in a separate loop. The resulting schedule is shown in Fig. 6.8(b).
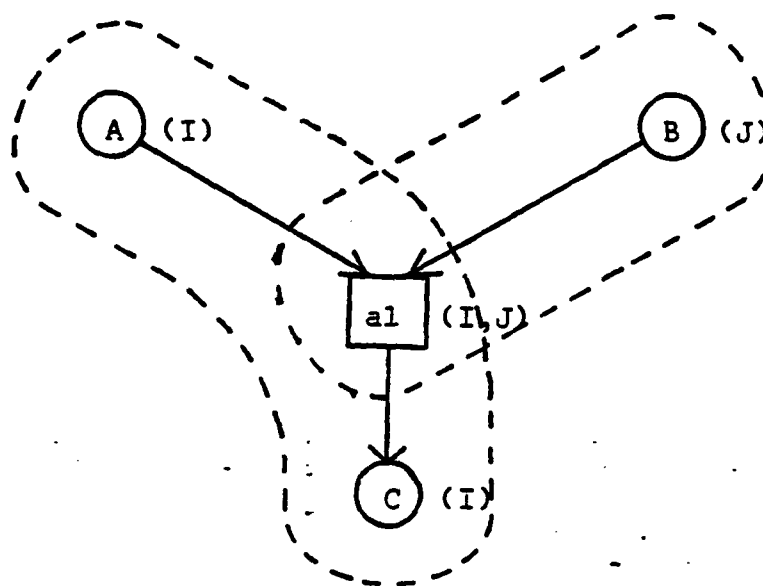
Fig. 6.8(a) An Array Graph to be scheduled
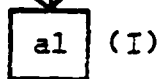
```
DO I ;
            A
END ;
DO J ;
            B
            al   (I)
END :
DO I ;
            C
END ;
```

Fig. 6.8(b) The outer level loop structure

6.6 THE SCHEDULING ALGORITHM

The scheduling algorithm, called SCHEDULE, is documented below. The overall process is illustrated in Fig. 6.9. The solid lines show procedure calls and the dashed lines show passing of parameters and returns. The SCHEDULE process starts with construction of a reduced form of the Array Graph, which will be modified in the course of scheduling and is also easier to manipulate. It then calls a recursive procedure SCHEDULE_GRAPH. This procedure accepts an Array Graph as input and returns a schedule as output. SCHEDULE_GRAPH calls on a number of procedures to perform its tasks. It calls first the procedure STRONG to construct a Component Graph out of the reduced Array Graph (or subgraphs of it in recursive calls).

Next, the major iteration in SCHEDULE_GRAPH schedules the outer loop scopes. This iteration repeats until all the components in the Component Graph have been scheduled. This major iteration loop finds first all the candidate ranges.

Next there is a nested iteration for trial scheduling of all the candidates ranges. It consists of calls to four procedures. Procedure INDRSUB is called first to find the range sets of each candidate range. If a candidate range has some subranges related to it, the sets of the subranges will also be included in the major range set. CLOSURE is then called to get the subgraph for the closure of the range

set. Then MAX_SCHED is called to do a trial scheduling. MAX_SCHED accepts as input a subgraph which consists of the closure of a respective range set and returns as output a loop scope which contains components in the closure of the range set that have been trial scheduled. The trial scheduling consists of repeated mergers into a loop scope of the components in the closure of the range set which do not depend on any other components. As a component is merged into the loop scope, it is deleted from the subgraph of closure of the range set. The merger repeats until no more components can be scheduled. Procedure EVALUATE is then called to compute the memory penalty associated with the loop scope.

At the end of the nested iterations for all the candidate ranges, SCHEDULE_GRAPH selects the loop scope with the smallest penalty. It will eventually form a part of the final schedule. The components in the selected loop scope are first merged into a single component and then marked off in the Component Graph.

The above major iteration loop is repeated, as noted above, until the Component Graph is empty. The outer loop scopes are thus all found. The corresponding components are topologically sorted. It is necessary then to find the nested loop scopes, if any, for each outer loop scope subgraph. As SCHEDULE_GRAPH selects the next component in the topological sorting, it calls the procedure EXTRACT to

extract these subgraphs, which correspond to the selected loop scopes. Each of these subgraphs must be internally scheduled. EXTRACT calls SCHEDULE_GRAPH recursively, to schedule each of the subgraphs. A component that is not within a loop scope needs not be further internally scheduled.
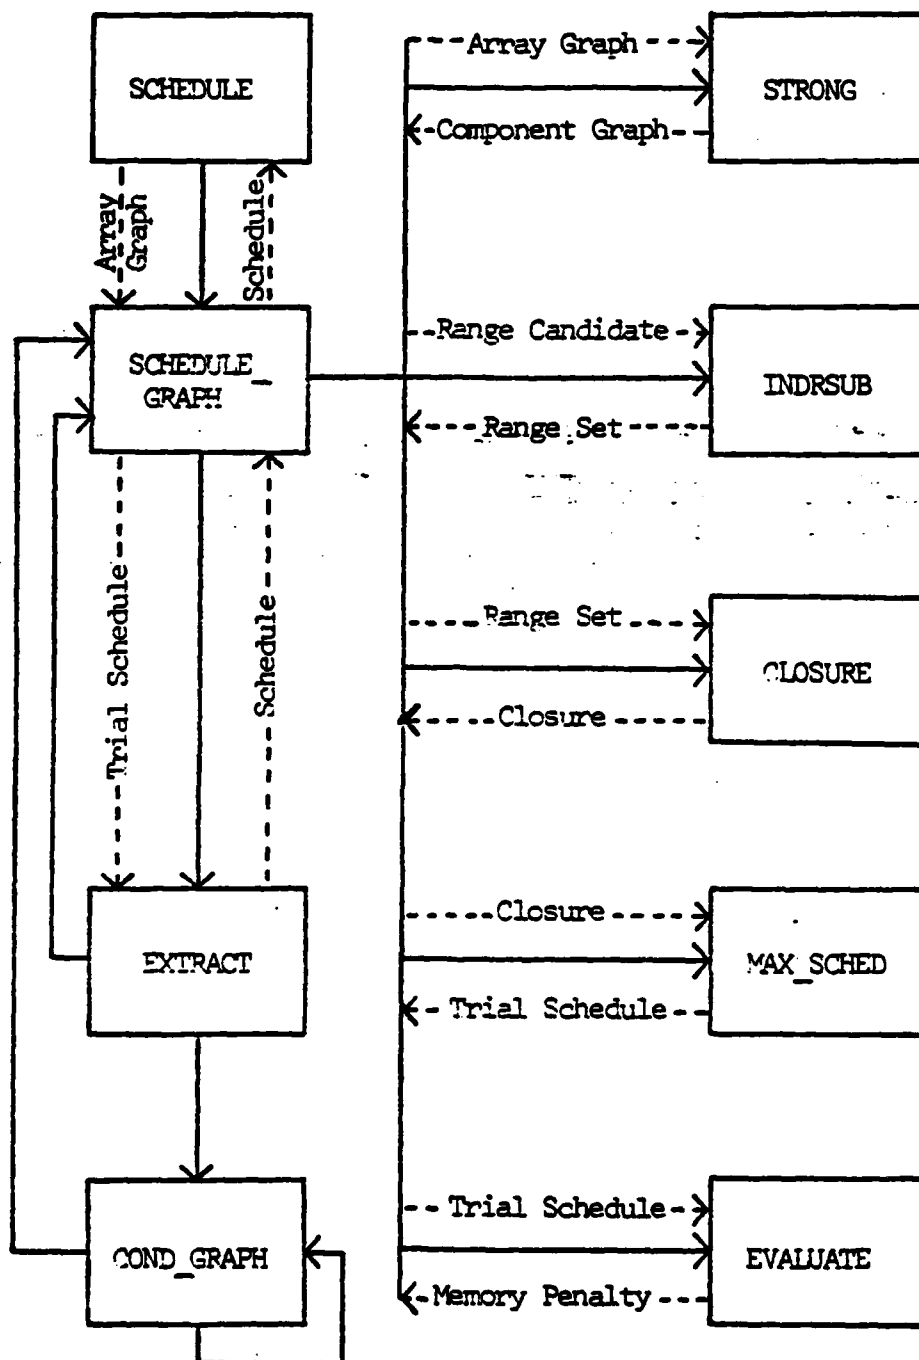
Fig. 6.9 Various components of the scheduling
algorithm

Global Data Structure for SCHEDULE

The reduced form Array Graph, constructed by the SCHEDULE procedure, consists of a list of elements of type GNODE, with the following fields:

NXT_GNODE - A pointer to the next element in the list. (At the generation of the reduced form Array Graph all the GNODEs form a single list. During the process separate lists will link the GNODEs in each MSCC.)

NODE_ID - The node number of the element in the dictionary.

SUXL - A pointer to a list of edges connecting this element to its successors. Initially this is identical to the SUCC_LIST list. As the process proceeds, some of the edges are removed from this list.

The components in the reduced Array Graph are found by the procedure STRONG. STRONG modifies the list connecting the nodes in the Array Graph to form separate lists for each MSCC.

The initial number of components in a Component Graph is denoted as COMP_CNT. Every component is assigned a component number from one to COMP_CNT. The component graph is defined in the following four vectors.

1) NODELST(COMP_CNT). Points to a list of GNODE elements in the Array Graph which belong to the respective component.

2) ACOMP(COMP_CNT). A boolean value showing whether the

component exists in the component graph or not. In the course of the process, when a component is merged into some other component, its corresponding ACOMP bit is reset.

3) INCMP(COMP_CNT). A boolean value showing whether a component has been scheduled or not. Once a component has been scheduled, its corresponding bit will be reset. Thereby it will not be scheduled again.

4) CEDGES(COMP_CNT). Points to a list of edges which originate from the component and end at its successor components. Every element in the list has two fields. One field contains the component number of its successor and the other is a pointer which points to the next edge.

A subgraph of the Component Grap. can be represented by a bit vector like INCMP. If a component is in the subgraph, its corresponding bit will be set. Otherwise, the corresponding bit will be reset. In the following, all the subgraphs of the Component Graph will use this representation.

The finally generated program schedule is structured as a list of schedule elements. There are four types of schedule elements: node-element, for-element, simul-element, and cond-element. A node-element corresponds to a primitive program event in the generated program such as the computation of an assertion, opening a file, reading a record. A for-element corresponds to a loop in the

program. The body of the loop is also represented by a schedule list and pointed to from the for-element. Similarly, a simul-element corresponds to an iterative computation for a simultaneous block and points to a list in the body of the iteration. The cond-element is used to represent a conditionally executed block which corresponds to the scope of a subrange. It will point to the respective body list.

1) A node-element is a structure NELMNT, with the following fields:

 NXT_NLMN – Pointer to the next element in the schedule.

 NLMN_TYPE – Equal to 1, denoting this is a node-element.

 NODE$ – The node number.

2) A for-element is a structure FELMNT, with the following fields:

 NXT_FLMN   – Pointer to the next element in the schedule.

 FLMN_TYPE – Equal to 2, denoting this is a for-element.

 ELMNT_LIST– Pointer to a program schedule which is the body of the loop.

 FOR_NAME   – The dictionary node number of the loop variable.

 FOR_RANGE – The dictionary node number where the range of the loop variable is specified.

3) A simul-element is a structure SELMNT which is used for a

simultaneous equation block. It has the same structure as FELMNT with FLMN_TYPE equal to 3.

4) A cond-element is used for a conditionally executed block. It has a similar data structure as FELMNT except that the field FLMN_TYPE is always equal to 4.

Algorithm 6.1 SCHEDULE_GRAPH

Input.

G: A pointer to the reduced Array Graph which is represented by a GNODE list.

L: The nesting level L.

Output.

A program schedule for the input graph G.

Data Structures.

GSIZE(COMP_CNT): The number of nodes in a component.

MINFREE(COMP_CNT): The minimum of the number of unscheduled dimensions associated with any node in a component.

SUBRNGR($RNG_SET,$RNG_SET): A boolean matrix which shows the subrange relationships. If the jth range set is a subrange of the ith range set, then SUBRNGR(i,j) will be set to '1'B.

RNG_VEC($RNG_SET): For each range set, it indicates the node number of the indirect indexing vector which reduces the major range into this range set, if any.

1. Call procedure STRONG to find out all the MSCCs in the

MICROCOPY RESOLUTION TEST CHART
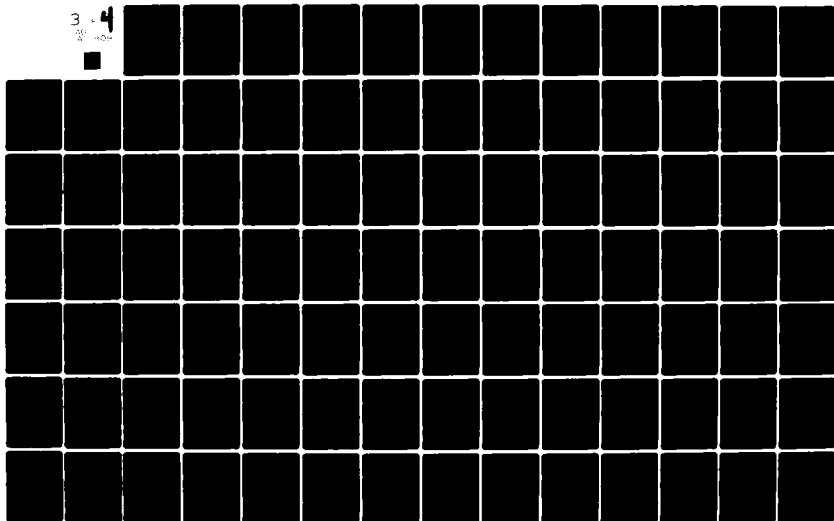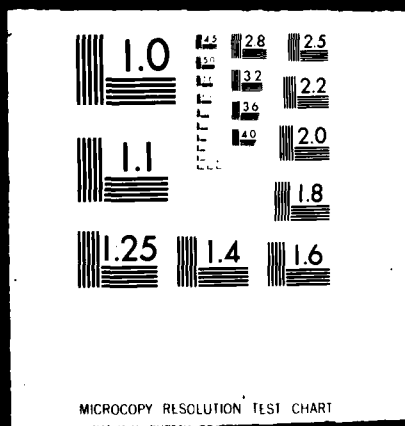
Array Graph G and then construct a Component Graph with each MSCC as a node. Initially all the components are put in the Component Graph and the corresponding ACOMP and INCMP bits are set to '1'B.

2. For each component, compute the corresponding element of the vector GSIZE, which is the number of nodes in the component, and the corresponding element in the vector MINFREE, which is the minimum of the number of unscheduled dimensions associated with any node in the component. Also compute the SUBRNGR matrix by scanning the indirect subscript expressions used in the assertions, and the vector RNG_VEC which gives for each range set number the node number of the indirect subscript, if any.

3. If a component has MINFREE=0, it is not to be scheduled in any loop. We will mark it off from the Component Graph by setting the corresponding INCMP bit to '0'B. This component will be a single component block.

4. Repeat step 5 to 11 to schedule all the outer level loops, until all components in the Component Graph have been marked off.

5. Select the ranges of node dimensions which are not yet scheduled and where the respective range does not have real arguments of unscheduled subscripts. The selected ranges can be scheduled in the outer level loops. The ranges of those node dimensions will be the candidate ranges.

6. Repeat step 7 to 10 for each range candidate. Steps 7 to 10 consist of a trial scheduling of a range candidate Ri.

7. Call procedure INDRSUB. This procedure computes a subgraph S which contains all the components which are in the range set of Ri or the range set of a subrange of Ri. S is represented as a bit map similar to INCMP.

8. Call procedure CLOSURE to find the subgraph S'=closure(S).

9. Call procedure MAX_SCHED with subgraph S' and range candidate Ri as input parameters to form a loop scope Li which contains a subgraph of S'. Li is represented as a bit map similar to INCMP.

10. Call procedure EVALUATE to compute the memory penalty of Li.

11. Choose the loop Lj with the smallest memory penalty. Merge all the components in Lj into one component, say Ck, by modifying the list pointed to by the NODELST of Ck to include all the GNODEs in the other merged components. ACOMP, INCMP, and CEDGES vectors are also modified to reflect the new component. Then set INCMP(k) to '0'B to mark the whole loop scope off from the Component Graph.

12. Do a topological sort over the resulting components of the component graph where each component corresponds to either a single node or a loop scope in the schedule to be returned.

13. Schedule each component separately. If there is no distinguished dimension for the nodes in a merged component, a node-element will be formed for the component. Otherwise, call the procedure EXTRACT to form a for-element for the component.

Algorithm 6.2 STRONG

Input.

G: A pointer to an Array Graph.

Output.

NODELST: A list of components which are the MSCCs of the input graph. Every component is represented by a list of GNODE elements which belong to the component.

1. Clear the stack, the component count, the list of components NODELST, and the variable COUNT. For each node v in the graph G set

   DFNUMBER(v) = 0

2. For each node v in the graph G such that DFNUMBER(v) = 0 call SEARCH(v) to add the components reachable from v to the component list NODELST.

3. Return the component list as the result.

Algorithm 6.3 SEARCH

Input.

v: A node in a graph which is not examined yet.

Output.

The NODELST for all the MSCCs reachable from node v.

1. Set COUNT to COUNT+1 and DFNUMBER(v), LOWLINK(v) to COUNT. Push v on the stack.

2. Repeat the following substeps for each node w, a direct descendant of v.

    2.1 If DFNUMBER(w)=0, call SEARCH(w) and then let LOWLINK(v)=min(LOWLINK(v),LOWLINK(w)).

    2.2. Else, if DFNUMBER(w)>0 and w is on the stack, then let LOWLINK(v)=min(DFNUMBER(w),LOWLINK(v)).

3. If LOWLINK(v)<DFNUMBER(v) then return.

4. Else, LOWLINK(v)=DFNUMBER(v). Node v is a root of a strongly connected component. All the elements (above and including v) on the stack are successively popped off the stack and linked into a list – a subgraph which is defined as a component. This component is placed on the top of a list of components pointed to by the variable COMP_LIST. In addition a unique component number is assigned to each node w in the current component.

Algorithm 6.4 INDRSUB(RANGE,GI)

Input.

RANGE: A candidate range (a range set number).

Output.

GI: A subgraph which contains all the components in the range set of RANGE and the components in the range sets of the subranges of RANGE which can be included

in the loop scope of RANGE.

1. Construct a subgraph GI which contains all the components in the Component Graph which have an unscheduled dimension with the range RANGE. GI is represented in a bit vector similar to INCMP. Set GI(k)='1'B if the kth component is in the range set of RANGE. The edges from these nodes are given in CEDGES.

2. If RANGE has no subranges, return GI as the result. This information stored previously in SUBRNGR matrix, which shows the subrange relationships.

3. Otherwise, repeat step 5 to 8 for each immediate subrange RNGIK of RANGE.

4. Call INDRSUB recursively with RNGIK as input parameter and GIK as the output parameter. GIK will contain the components which can be scheduled in the loop of RNGIK.

5. Call procedure CLOSURE to compute the closure of GIK in the Component Graph. Then put the closure into GIK.

6. Set the union of GI and GIK into GI. (Note that this may be reversed in step 8.)

7. Call MAX_SCHED procedure to do a trial scheduling for subgraph GI.

8. If the subgrpah GI can not be scheduled completely, then at least one node, and possibly more, will have to be physical. Also the range specification of the subrange may become necessary. Therefore we decided that in this case it is not worthwhile to merge the range set of RNGIK with the range set of RANGE and GIK is taken out

of GI.

9. Return GI as the result.

Algorithm 6.5 CLOSURE(COMPS)

Input.

COMPS(COMP_CNT): A bit vector with a set of components marked by '1'B. Other components are marked by '0'B.

The algorithm also uses the global data structures (ACOMP and CEDGES).

Output.

CCOMPS: A bit vector with the closure of the set of components in the input marked by '1'B. Other components are marked by '0'B.

1. Create a bit vector NACOMP (size COMP_CNT) with the components in ACOMP marked except the components in COMPS are merged into one component. This also involves creating a vector NCEDGES similar to CEDGES except reflecting the merger of the components in COMPS.

2. Find all the MSCCs in the new component graph (consisting of the new vectors NACOMP and NCEDGES).

3. Locate the MSCC which includes the components in COMPS.

4. Construct CCOMPS, a bit vector (size COMP_CNT), with all the components in the MSCC marked. This is the closure set of the input.

Algorithm 6.6 MAX_SCHED

Input.

INCMP: A bit vector where a set of yet unscheduled components is marked by '1'B. Other scheduled components have a value '0'B. Note that these unscheduled components are the basic MSCCs found by STRONG. The function of MAX_SCHED is to schedule as many of the marked components as possible.

MERGCMP: A bit vector with the closure of a range set marked by '1'B.

RANGE: The candidate range (range set number).

Output.

COMPS: A bit vector with the components, which have been trial scheduled in a loop, marked by '1'B.

POSITION: A vector (size is DICTIND- the number of nodes in the dictionary). The position in each scheduled node of the distinguished dimensions that corresponds to the loop parameter.

1. Initialize the POSITION entries to 0.

2. For each component i, if INCMP(i)='1'B (i.e. it is not yet scheduled), MERGCMP(i)='1'B (i.e. it is in the closure set), then search the CEDGES vector and set PREDCNT(i) to number of predecessors in MERGCMP. If PREDCNT(i)=0 then put component i into a list of candidates to be trial scheduled.

3. Repeat steps 4 to 8 until the list (referred to in step 2) is empty. The function of steps 4 to 8 is to merge one component from the list into the loop scope

represented by COMPS.

4. Remove a component, say C_i, from the list. Search through the NODELST of C_i, if there exists a node v with POSITION(v)>0 (i.e. its distinguished dimension has been determined in a previous iteration), then set FIRSTNODE=v, and go to step 7.

5. Else, arbitrarily pick any node of the component. Let it be denoted by v. Set FIRSTNODE=v.

6. Search the subscript list of node v until finding a dimension j that has not been scheduled in a loop scope (i.e. IDWITH=0) and its range is the same as the RANGE parameter. If found, then POSITION(v)=j. If none found then this component should not be scheduled in the loop scope. Therefore go to next iteration (i.e. end of step 9).

7. Propagate the distinguished dimension of node v repeatly until all the nodes in C_i have their distinguished dimensions defined. During each propagation step:

    7.1 Propagate the distinguished dimension forward along the edges originated from node v to all the nodes at the terminating end of the edges.

    7.2 If the node to which a distinguished dimension is propagated does not belong to C_i then do not further propagating the distinguished dimension from this node forwards.

    7.3 If propagation is not possible to any node in C_i because of type 4 subscript expression then the

current iteration may be terminated, i.e. go to end of step 9.

8. The current component can be merged into the loop scope. Set COMPS(i)='1'B.

9. Search through the list pointed by CEDGES(i). For every edge from Ci to Ck set PREDCNT(k)=PREDCNT(k)-1. If PREDCNT(k)=0, INCMP(k)='1'B, and MERGCMP(k)='1'B, then put Ck into candidate list.

Algorithm 6.7 EVALUATE

Function: Given a loop scope, compute the resulting penalty in use of memory. This procedure is called after each trial schedule for a range candidate and again after the final schedule was selected.

Input.

COMPS: A bit vector of size COMP_CNT with the bits correspondning to components in a loop scope equal to '1'B.

EVAL_SET: A bit denoting whether EVALUATE is called to evaluate memory penalty of a trial schedule or for the selected schedule, in which case the selected memory allocations are recorded in STOTYP.

Output.

PENALTY: The memory penalty of the loop scope, in bytes.

Data structure.

SRCPHY, TGTPHY: When an edge in an Array Graph crosses a

boundary of a loop scope then, depending on the type of the edge, the memory allocation for the data node at the origin or terminating ends of the edge may have to be physical. The SRCPHY bit vector denotes for each type of edge (there are 28 types) whether the memory allocated to the node at the origin end of the edge (the source node) must be physical. Similarly, the TGTPHY vector refers to the node at the terminating end of the edge (the target node).

MRAL: The memory requirement, in bytes, after the loop is formed.

MRIC: The memory requirement in the ideal case.

STOTYP: A field in the data structure LOCAL_SUB. For a virtual dimension, STOTYP=0. For a window of width k+1 dimension, STOTYP=k+1. For a physical dimension with upper bound u, STOTYP=-u.

1. Repeat steps 2 to 6 for every edge in the Array Graph. Each iteration computes the effect of the edge on use of memory.

2. If the source and the target nodes of the edge are in COMPS, this is an internal edge, then go to step 6 to examine the subscript expression of the edge to determine its effect on use of memory.

3. If both the source and the target nodes of the edge are not in COMPS, then this edge has no effect on memory useage. Go to end of iteration, at end of step 6.

4. If none of the above then the edge crosses the loop

boundary. In this case, if SRCPHY(EDGE_TYPE)=1, then the distinguished dimension of the source node must be physical. If TGTPHY(EDGE_TYPE)=1, then the distinguished dimension of the target node must be physical. The respective node numbers and the requirements for physical memory allocation are stored in a list. Also in this case go to the end of the iteration (at end of step 5).

5. If the subscript expression is of the form I-k and SRCPHY(EDGE_TYPE)=1, then the memory allocation for the distinguished dimension of the source node must be a window of width k+1. This is also stored in the list.

6. PENALTY is initialized to zero.

7. Repeat steps 8 to 11 for every node in the above list. These nodes have either a physical or window of width k+1 memory allocation. An iteration computes the memory requirement for a respective node.

8. In the case of a physical distinguished dimension, compute MRAL, as the product of all the ranges of the unscheduled node subscripts. In the case of a window of width k+1 for the distinguished dimension, compute MRAL as the product of k+1 and the ranges of the other unscheduled node subscripts.

9. To compute MRIC it is necessary to scan each unscheduled node subscript. If its storage type STOTYP is 0, then the ideal memory requirement for this dimension is one. If STOTYP<0, the memory allocation has previously been

determined as physical, then the ideal memory requirement is -STOTYP (u). MRIC is the product of these ideal ranges.

10. The penalty for the array node ND_PENALTY= (MRAL-MRIC)*(length of node element in bytes).

11. PENALTY=PENALTY+ND_PENALTY.

12. If EVAL_SET='1'B then if the distinguished dimension is physical then STOTYP in every unscheduled dimension is equal to the minus of its range, if the distinguished dimension is a window of width k+1 then STOTYP of the distinguished dimension is k+1 and for the other unscheduled dimensions STOTYP is the minus of their respective range.

Algorithm 6.8 EXTRACT

Function: To obtain the for-element for a loop, including the schedule elements for the body of the loop scope.

Input.

SUBGRAPH: A pointer to a reduced Array Graph of the component scheduled into one loop scope.

SVPOSITION: A vector with an element for every node in the SUBGRAPH. Each element has the value of the dimension number of the distinguished dimension of the respective node.

L : The nesting level.

Output.

A for-element which is the schedule of the input graph.

1. Allocate a for-element. Set FOR_NAME to loop parameter name and FOR_RANGE to the range set number of the loop parameter.

2. If the current loop range has some immediate subranges, then call procedure COND_GRAPH and upon return go to step 7. COND_GRAPH takes over all further scheduling of a body of a loop which contains conditionally executable nodes due to use of indirect subscripting.

3. Delete all the edges from the graph with distinguished dimension subscript expressions of type 2 or 3. The precedence expressed by these edges is assured by the order of the iterations.

4. Set IDWITH of the distinguished dimension of all the nodes in the subgraph to L, the nesting level of the current loop.

5. Call SCHEDULE_GRAPH, with SUBGRAPH and L+1 as the parameters, to get the schedule of the resulting graph.

6. Set ELMNT_LIST in the for-element structure to point to the schedule returned from step 5.

7. Return the for-element as output.

<u>Algorithm</u> <u>6.9</u> COND_GRAPH(TOP_RANGE,GRAPH)

Function : To obtain the schedule elements of the body of a loop scope, which includes cond-elements.

Input.

TOP_RANGE: The range set number of the highest level major range in the SGRAPH.

SGRAPH: A graph to be scheduled within an iteration block of the range TOP_RANGE.

Output. A schedule for SGRAPH.

1. Scan all edges in SGRAPH. If an edge has a subscript expression in the distinguished dimension of types 2, 3, 6, or 7, and either the source or the target nodes have the TOP_RANGE range, then delete this edge from SGRAPH.

2. If node X is the indirect indexing vector served to reduce the range TOP_RANGE to a subrange RNGIK, then draw an edge from X to all the nodes in the range set of RNGIK.

3. Call procedure STRONG to form a Component Graph for SGRAPH, consisting of ACOMP and INCMP, CEDGES, and NODELST. ACOMP and INCMP are bit vectors (the size is the number of MSCC found by STRONG). These vectors are all of the value '1'B.

4. For every subrange RNGIK of TOP_RANGE, merge all the components in the range sets of RNGIK or its direct and indirect subranges into one component. Set the INCMP vector elements of the merged components to '0'B.

5. Repeat steps 6 to 9 until all the elements in INCMP are '0'B. Each iteration merges a group of components with TOP_RANGE range.

6. Call CLOSURE with INCMP to obtain the closure set MERGE_CMP.

7. CALL MAX_SCHED with INCMP, MERGE_CMP, and TOP_RANGE. It returns CCOMPS.

8. Merge the components in CCOMPS into one component, updating NODELST, CEDGES, ACOMP, and INCMP.

9. Set the element of INCMP corresponding to the merged schedule to '0'B.

10. Repeat steps 12 to 13 for the components in ACOMP.

11. Select the next component in ACOMP in a topologically sorted order. Let this component be COMPI.

12. Let RNGIK be the range of the component COMPI. If RNGIK=TOP_RANGE, then mark the distinguished dimension of each node in the component as scheduled and call procedure SCHEDULE_GRAPH to get a schedule for this component. Go to step 14.

13. Otherwise, allocate a cond-element to this component. Call procedure COND_GRAPH recursively with RNGIK and COMPI as the input parameters to get a schedule for the conditional element.

14. Return the schedule elements obtained as the final schedule of SGRAPH. Note that the order of the schedule elements was determined by the selection of components in a topologically sorted order in step 11. The schedule elements are obtained either in step 12 or 13, depending on whether they are cond-elements or other elements respectively.

CHAPTER 7

CODE GENERATION

7.1 OVERVIEW OF THE CODE GENERATION PROCESS

Code Generation is the last phase of the processor. It uses the data structure generated in Array Graph construction, specification analysis, and program scheduling. As shown in Fig. 7.1 the code generation process accepts two inputs: the program schedule created in the scheduling phase and attribute tables produced in the analysis phase. Recall that the program schedule is an ordered sequence of schedule elements described in section 6.6. The nodes referenced in schedule elements can be found in the dictionary. The attributes of the respective nodes are in the dictionary. They are described in the section 4.2.1. The output is a complete PL/I program ready for compilation. The executable PL/I code is written out to the "PL1EX" file. The PL/I "ON" conditions are written to the "PL1ON" file and the PL/I code for declaring the object data items is written to a "PL1DCL" file.
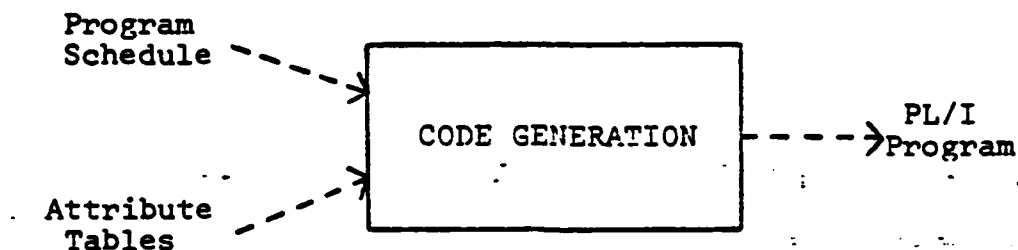
Fig. 7.1 Overview of the Code Generation Phase

Fig. 7.2 shows the overall organization of the code generation process, consisting of the main procedure CODEGEN which in turn calls on the other procedures to perform certain tasks. The PL/I execution code is generated by the GENERATE procedure which examines the elements of the schedule one at a time, and invokes the procedures that are indicated by types of program events. The GPL1DCL procedure generates the data declarations. GENERATE calls GEN_NODE to generate statement for node elements of the schedule. The GEN_NODE calls on GENIOCD for input-output operations and on GENASSR for assertions. GENERATE also calls GENDO and GENEND for generating iteration control structures for for-elements, and on COND_BLK and COND_END for generating conditional block statements for cond-elements. These procedures are briefly reviewed in section 7.2. They are described in greater detail together with other auxiliary tasks in the subsequent sections that follow.
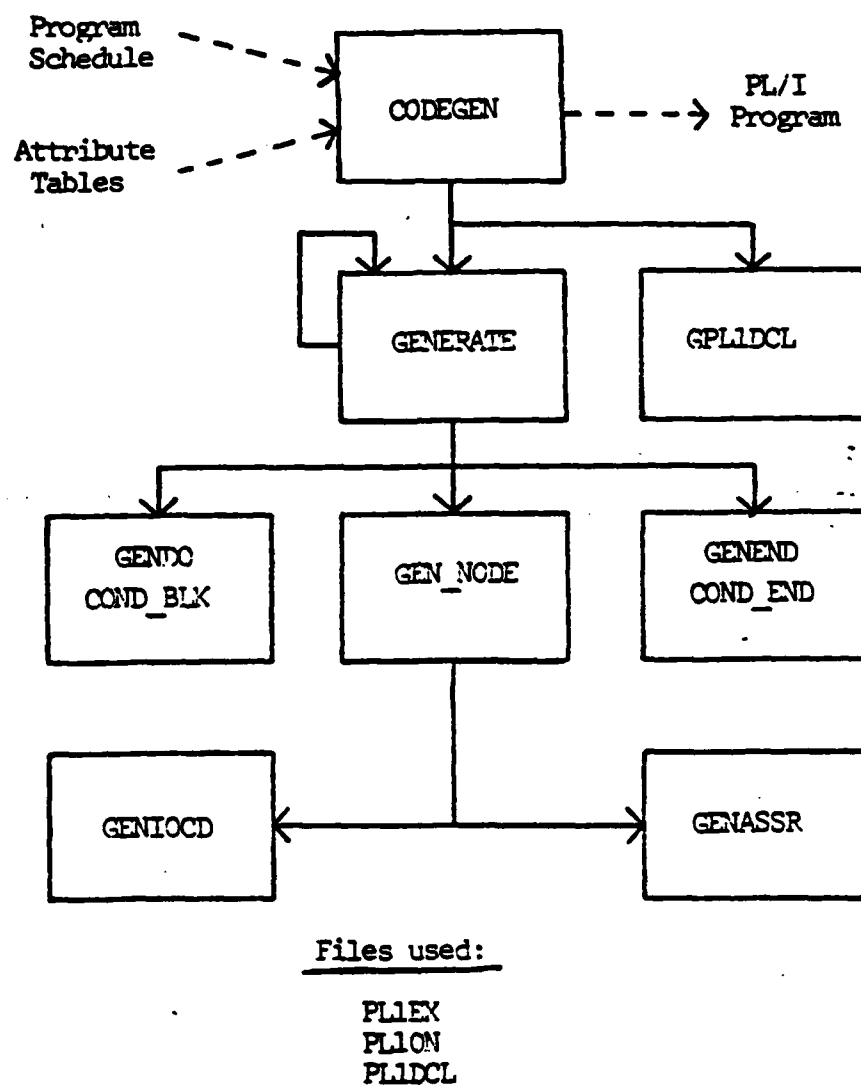
Fig. 7.2 Components of Generating PL/I Code

7.2   THE MAJOR PROCEDURES FOR CODE GENERATION

7.2.1 CODEGEN - THE MAIN PROCEDURE

CODEGEN starts with opening the output files PL1EX,
PL1ON, and PL1DCL. It next generates code that will handle
program errors. Most of these errors are due to input data
errors discovered by data type conversions in the program.
The user can also define additional error conditions. The
statements written to the PL1EX file are as follows:

```
ALLOCATE ERROR, ACC_ERROR ;

ACC_ERROR = '0'B ;

ALLOCATE $ERR_LAB ;

$ERR_LAB = END_PROGRAM ;
```

The declarations written to the PL1DCL file are as follows:

```
DCL (ERROR, ACC_ERR, NOT_DONE) CTL BIT(1) ;

DCL $ERR_LAB LABEL CTL ;
```

Finally the ON condition code is sent to the PL1ON file as
follows:

```
ON ERROR

BEGIN

  /* write erronous input record to ERRORF file */
  WRITE FILE(ERRORF) FROM($ERROR_BUF) ;
  ERROR = '1'B ;      /* set error flag           */
  GO TO $ERR_LAB ;   /* go to end of loop where */
END ;                 /* error was detected       */

ERROR_RESTART:
```

CODEGEN next passes the entire program schedule to GENERATE, which will generate the portions of the program for the schedule elements. When this is completed CODEGEN passes the attribute tables to GPL1DCL to generate data declarations. Finally CODEGEN calls on MERGEPL1 to merge the three output files.

7.2.2 GENERATE - INTERPRETING SCHEDULE ELEMENTS

This recursive procedure scans the schedule given by the list of schedule elements, LIST, for a loop nesting level LEVEL. To start with, CODEGEN passes the whole schedule at level 0. In subsequent calls GENERATE will receive a schedule of a loop scope at each nesting level. GENERATE calls lower level procedures to process the different types of schedule elements as follows:

1. Scan each element of the list LIST. For each element perform steps 2 to 4.

2. If the element is a node-element call GEN_NODE which will generate the code for the schedule element.

3. If the element is a for-element do the following:

    3.1 Call GENDO to produce a code for opening a loop.

    3.2 Call GENERATE recursively with the list of the elements within the loop's scope and level = LEVEL+1.

    3.3 Call GENEND to generate the termination of the loop.

4. If the element is a cond-element do the following:

4.1 Call COND_BLK to produce the code for opening a conditional block.

4.2 Call GENERATE recursively with the list of the elements within the condition block and level = LEVEL.

4.3 Call COND_END to generate the termination of the conditional block.

7.2.3 GENDO - TO INITIATE THE SCOPE OF ITERATIONS

This procedure produces the code for a control statement initiating an iteration loop. The loop variable name FORNAME and the termination criterion are taken from the fields FOR_NAME and FOR_RANGE in the for-element being scanned.

The following instructions are intended for recovery from a program error. They always precede each loop control statement:

```
ALLOCATE ERROR, ACC_ERROR ;
/* reset accumulative error flag */
ACC_ERROR = '0'B ;
ALLOCATE $ERR_LAB ;
$ERR_LAB = LOOP_ENDc ;
```

The "c" following LOOP_END is a unique number assigned to the loop. The purpose of these statements is to ensure that an error occurring within the loop scope will cause the

control be directed to LOOP_ENDc which is a label immediately preceeding the end of the loop.

The DO-statement itself is constructed next. Two basic forms for the loop control statements are used:

1)

        DO name = 1 TO upper [WHILE (condition)] ;

2)

        name = 0 ;

        DO WHILE (condition) ;

          name = name+1 ;

"name" is the loop variable. "condition" is the termination condition.

If the termination criterion given is that of a fixed upper limit or given through a SIZE variable, the first form is used and "upper" is either a constant number or a variable of the form SIZE$X.

If the range is specified by an END.X control variable, the second form of loop control is used. In this case we use NOT_DONE in the condition and the following statements are generated before the beginning of the loop:

        ALLOCATE NOT_DONE ;

        NOT_DONE = '1'B ;

NOT_DONE will be reset to '0'B whenever the appropriate END.X variable is set to 'true'.

If there is an end-of-file condition associated with the iteration, either as the main termination condition, or because this is an iteration on an input record or group above the record level which are last in their peer group, we add:

        ^ENDFILE$file

to the condition "condition".

7.2.4 GENEND - TO TERMINATE THE SCOPE OF ITERATIONS

This procedure produces the code needed at the end of the loop scope. Since at times, we use k+1 locations to store a window of size k+1 of an array, it is necessary on each iteration to shift the window by one element position. This is done at the end of the iteration. The size of respective window is originally stored in STOTYP of the node subscript of each array node. GENERATE passes the node numbers of arrays using window dimensions in a list called PREDLIST to GEN_END. Based on this list GEN_END generates statements to shift the window by one element position. The actual range declared for a window dimension is k+1. In each iteration we compute (or read) $A(\ldots, k+1, \ldots)$ and may refer to the previous element as $A(\ldots, k, \ldots)$. When an iteration is completed we transfer $A(\ldots, I+1,\ldots)$ to $A(\ldots, I,\ldots)$ for I from 1 to k.

After producing a sequence of these shifting operations we produce the label:

    LOOP_ENDc: ;

where "c" is the unique count associated with the current loop. If the termination criterion for the loop was through an END.X control variable we also produce the code:

    IF END.X THEN NOT_DONE = '0'B ;

This has to be done at the end of the loop since the value of END.X at a given iteration determines whether this iteration will be the last.

After this we produce the following statements:

    $TMP_ERROR = ACC_ERROR ;

    FREE ERROR, ACC_ERROR ;

    FREE $ERR_LAB ;

    IF $TMP_ERROR THEN ERROR, ACC_ERROR = '1'B ;

If the termination criterion was through an END.X control variable we also produce:

    FREE NOT_DONE ;

7.2.5 COND_BLK - INITIATE A CONDITIONAL BLOCK

This procedure produces the code necessary to initiate a conditional block. The conditional block will be executed within the iteration only when the value of tne indirect subscript is increased. The indirect subscript node number

is stored in the FOR_RANGE field of the cond-element being scanned. An IF-statement is generated to test the above condition. Inside the conditional block we will use a new symbol for the indirect subscript. For example, if X(I) is the indirect subscript then we define a new subscript J=X(I). Let 'old-sub' denote the subscript running in the major range, i.e. I. The 'new-sub' denotes the new representation of the indirect subscript, i.e. J. A boolean variable, $B_X, indicates whether the conditional block should be executed. The code to compute $B_X is generated by GEN_NODE when the node X is scanned in the schedule. The new-sub is of the form $Xn where 'n' is a unique number associated with this conditional block. The following declaration statements are issued:

```
DCL $Xn FIXED BIN ;
DCL $B_X BIT(1) ;
```

The following codes is then produced:

```
IF $B_X THEN DO ;
    new-sub = X( ..., old-sub) ;
```

7.2.6 COND_END - TERMINATE A CONDITIONAL BLOCK

This procedure produces the code at the end of a conditional block. The above IF-statement has been generated by COND_BLK. Here we issue an 'END' statement to terminate the IF-statement.

7.3 GEN_NODE - CODE GENERATION FOR A NODE

This procedure generates the code associated with a schedule node-element. It branches to different parts according to the types of nodes.

7.3.1 PROGRAM HEADING

If the node is a module name (type MODL) we produce the code:

        name: PROCEDURE OPTIONS(MAIN) ;

This code is routed to the file PL1DCL.

7.3.2 FILES

If the node is a file node (type FILE) we first generate three names. "file_stem" is the file name with prefixes "NEW" or "OLD" removed, if any. "name" is the full name of the node, including all prefixes. "file_suff" is the file_stem with the suffix of 'S' for source file, 'T' for target file, and 'U' for update file (both source and target). The following declaration statements are routed to PL1DCL file.

        DCL name_S CHAR(length) VARYING INIT(' ') ;
        DCL name_INDX FIXED BIN ;

"length" is the maxmimum length of records in the file.

"name_S" is the name of a buffer into which records in the file are read. (It is VARYING as the file may have more than one record type, with different lengths.) "name_INDX" is a variable used to scan the buffer for packing and unpacking the records (explained further later).

1. If the file is an input file we produce the statement:

        OPEN FILE (file_suff) ;

2. If the file is a sequential input file and an end-of-file is not explicitly mentioned by the user, we produce the declarations:

        DCL ENDFILE$file_stem BIT(1) INIT('0'B) ;
        DCL $FSTfile_suff BIT(1) INIT('1'B) ;

   routed to PL1DCL file. If the user explicitly mentioned the end-of-file variable then these statements will be generated when the declaration are generated for all variables by GPL1DCL.

        The statements:
          ON ENDFILE (file_suff)
          BEGIN
            ENDFILE$file_stem = '1'B ;
            name_S = COPY(' ',length) ;
          END ;

   are sent to PL1ON file. The purpose of these statements is to have the file buffer filled with blank characters when an end of file condition occurs.

3. If the file is an output file we produce the statement:

```
        CLOSE FILE(file_suff) ;
```

7.3.3 RECORDS

If the node is a record (type RECD) we call GENIOCD to produce the code for the reading or writing of records.

7.3.4 FIELDS

To process fields GEN_NODE calls procedure GENITEM. GEN_NODE also calls CHECK_VIRT to find if the node has a windowed dimension.   If the field node is an indirect subscript, X, the following code is issued.

```
        IF loop_var=1 THEN DO ;

         bname = '1'B; rname = 0; END ;

        ELSE IF X(loop_var)>X(loop_var-1) THEN DO ;

          bname = '1'B; rname = 0; END ;

        ELSE DO ;

          bname = '0'B; rname = 1; END ;
```

where loop_var is the current level loop variable, bname is of the form $B_X, and rname is of the form $R_X. Recall that bname indicates whether the associated conditional block will be executed. rname will be used to compute the index to reference an element such as A(X(loop-var)) in the case that array A has a windowed dimension.   This is explained further later in connection with the code

generation for assertions.

7.3.5 ASSERTIONS

If the node is an assertion we call the procedure GENASSR to produce the code for an assertion.

7.4 GENASSR - GENERATING CODE FOR ASSERTIONS

This procedure generates code for assertions. The main task of GENASSR is to transform the syntax tree representation of the assertion into a string representation acceptable by the PL/I compiler. The transformation is carried out by a recursive climb on the syntax tree, combining for each node the string representations of the descendant subtrees into a string representation of the tree rooted at that node. However, before performing the main task the procedure transforms assertions containing conditional expressions into conditional assertions. Thus, an assertion of the form:

```
        Y = IF (IF X>0 THEN Y>0 ELSE Y<=0) THEN X*Y

                                        ELSE -X*Y ;
```

will be transformed into:

```
        IF X>0 THEN IF Y>0  THEN Y = X*Y ;

                            ELSE Y = -X*Y ;

                ELSE IF Y<=0 THEN Y = X*Y ;
```

$$ELSE\ Y = -X*Y\ ;$$

The overall execution of GENASSR can therefore be summarily described as:

1. Transform assertions with conditional expressions into conditional assertions.

2. Form the string representation of the assertion.

7.4.1 TRANSFORMING CONDITIONAL EXPRESSIONS

This task is carried out by the procedure SCAN which uses the auxiliary procedure EXTRACT_COND.

7.4.1.1 SCAN (IN)

The procedure SCAN effects the complete transformation of assertions containing conditional expressions into conditional assertions. The procedure is presented with an assertion pointed to by IN, and returns a pointer to the transformed assertion. The steps in this procedure are as follows:

1. Check the root of the tree pointed to by IN to see whether it is a simple assertion or a conditional assertion. If it is a simple assertion then go to step 5.

2. We check next if the conditional assertion contains

conditional expressions. A conditional assertion has the form:

IF COND THEN S1 ELSE S2

where S1, S2 are assertions.

SCAN calls EXTRACT_COND to check whether COND contains a conditional expression. If COND contains a conditional expression, then EXTRACT_COND returns C, L, and R which are the parts of COND as follows:

COND = IF C THEN L ELSE R.

Otherwise, go to step 4.

3. If a conditional expression is found in COND then:

   3.1 SCAN then transforms the tree (pointed to by IN) into a tree IN1 which consists of the form:

   IF C THEN IF L THEN S1

   ELSE S2

   ELSE IF R THEN S1

   ELSE S2

   3.2 SCAN calls SCAN(IN1) recursively to further search for conditional expressions in IN1 and return a transformed conditional assertion.

   3.3 The transformed assertion is returned by SCAN.

4. If COND does not contain embedded conditional expressions, then there are two recursive calls to SCAN for the assertions S1 and S2 in IN. SCAN then returns the following assertion and exits.

   IF COND THEN SCAN(S1) ELSE SCAN(S2)

5. In the case of a simple assertion:

Y = E.

SCAN calls EXTRACT_COND(E) to search for conditional expressions in E. If none found, then assertion Y = E is returned unchanged. Otherwise, EXTRACT_COND returns C, L, and R which are the parts of E as follows:

E = IF C THEN L ELSE R.

6. If E contains conditional expression, then SCAN calls SCAN(IN2) recursively, where IN2 points to a tree of an expression of the form:

'IF C THEN Y = L

ELSE Y = R'

The return from the recursive call on SCAN is returned by SCAN as the transformed assertion.

7.4.1.2 EXTRACT_COND(ROOT,COND,LEFT,RIGHT)

This procedure identifies and extracts the leftmost conditional expression in a given expression pointed to by ROOT.

If a conditional expression is found the (pointer to the) condition is returned in COND and its first (THEN) and second (ELSE) subexpressions returned in LEFT and RIGHT respectively. If the analyzed expression contains no conditional expression the procedure returns NULL in COND.

Its operation is as follows:

1. Inspect the top level node of the given syntax tree.

2. If it is a conditional expression, return respectively the condition, the subexpression following THEN, and the subexpression following ELSE, then exit.

3. If the expression is a simple expression, i.e. a constant or a variable, return NULL and exit.

4. If the expression is a compound expression, scan each of its descendants by calling EXTRACT_COND recursively. Consider the first COND, LEFT, and RIGHT which are returned such that COND is not equal to NULL. In general, a compound expression is of the form:

$$E = g(E1,...,Em)$$

Assume that the recursive scanning of E1, ..., Em produces first COND not equal to NULL for Ei where 1<=i<=m, returning also the THEN and ELSE subexpressions L, and R respectively. Then the current call for E returns:

COND as the condition,

$g(E1, ...,Ei-1,L, ...,Em)$ as LEFT, and

$g(E1, ...,Ei-1,R, ...,Em)$ as RIGHT.

Thus the overall effect of EXTRACT_COND on an expression E is to extract a condition C if one exists in E (returned as COND), and then to compute E1 when C is true, and E2 when C is false. E1 and E2 are returned in LEFT and RIGHT respectively. Described in another way we look for C, E1, and E2 such that the following equivalence holds:

$$E = \text{IF C THEN E1 ELSE E2 .}$$

In particular this gives:

$$g(\text{E1, } \ldots, \text{Ei-1,(IF C THEN L ELSE R),} \ldots \text{Em}) =$$

$$\text{IF C THEN } g(\text{E1, } \ldots, \text{Ei-1,L,} \ldots, \text{Em})$$

$$\text{ELSE } g(\text{E1, } \ldots, \text{Ei-1,R,} \ldots, \text{Em}).$$

7.4.2 PRINT - TRANSFORMING THE ASSERTION INTO STRING FORM

This procedure is presented with a pointer to an assertion syntax tree and it converts the assertion tree into a string representation.

The procedure branches according to the types of the nodes in the assertion tree.

1. If the node is a subscripted variable A(E1,...,Em) we generate the string 'A('. We then scan each of the subscript expression E1 to Em and add them to the string according to the following subcases:

   1.1 If the dimension at position i corresponds to the dimension declared for repetition of a record and the variable A includes the prefixed 'NEXT', then

      1.1.1 If the dimension is scheduled as a window of width k+1 we insert the subscript value k+2.

      1.1.2 If the dimension is scheduled as physical and the expression Ei is a constant c, then insert the value of c+1. (See further below.)

      1.1.3 If the dimension is scheduled as physical and

Ei is an expression we call PRINT(Ei) and insert the returned value concatenated with '+1'.

1.2 If the dimension at position i is scheduled as a window of width k+1, in this case the physical allocation for the array dimension is k+2 elements with the k+1th element standing for the current value and the k+2th element standing for the field in the next record. The different subscript expressions are handled as follows:

1.2.1 If it is a simple subscript then we insert an integer k+1 as the subscript.

1.2.2 If the subscript expression is I-c, then an integer k+1-c is inserted.

1.2.3 If the subscript expression is X(I), then k+1-$R_X is inserted where k+1-$R_X points to the element A(X(I)). If X(I)=X(I-1) then $R_X is equal to 1, and if X(I)>X(I-1) then $R_X is equal to 0. (The code to compute $R_X is generated by GEN_NODE right after node X is scanned.)

1.2.4 If the subscript expression is X(I)-c, then k+1-$R_X-c is inserted as subscript.

1.2.5 If the subscript expression is X(I-a), then k-[X(I-1)-X(I-a)] is inserted as the subscript. X(I-1)-X(I-a) is the offset of A(X(I-a)) to A(X(I-1)) which is stored in the kth element of

the window for the ith dimension of array A.

1.2.6 If the subscript expression is $X(I-a)-c$, then $k-[X(I-1)-X(I-a)]-c$ is inserted as the subscript.

1.3 If the ith dimension of array A is physical and Ei is the subscript expression, we call PRINT(Ei) and insert the returned value.

2. For all other compound nodes we call PRINT recursively to convert the descendants and insert between them the string representation of the separators, operators, and delimiters. The latters are stored in the OP_CODE fields as integer codes. The integer codes are translated into the operator representation using the array KEYS and then inserted.

3. For atomic nodes we use the variable name either directly or through its node number. Loop variables (subscripts) are accessed through the level indication available in their IDWITH field which is used as an index to the array LOOP_VARS. Function names are retrieved by their function number indexing the table FCNAMES.

7.5 GENIOCD - GENERATING INPUT/OUTPUT CODE

GENIOCD is invoked by CODEGEN upon scanning a schedule element which corresponds to a record node. It accepts as input the node number in the schedule element. GENIOCD

generates PL/I READ, WRITE, or REWRITE statements with the appropriate parameters, based on the attributes of the file, as well as the control code or condition code associated with the input/output operation.

Table 7.1 summarizes the different statements generated by GENIOCD for the different cases. Each of the different cases in Table 7.1 shows the conditions defining the case and the statements which are generated for the case. The upper case letters represent the part of the actual PL/I string being generated, whereas the lower case letters are the metanames of the items obtained from the program schedule elements.

Several preparatory steps are taken before branching to the different cases.

1. Definition of names: We generate several variable names derived from the record name that will be used in the code. Let the record name be designated by rec.

   1.1 If rec is of the form OLD.X or NEW.X we define recname as OLD_X or NEW_X respectively.

   1.2 Otherwise we define recname as rec.

   1.3 Recbuf is defined as recname_S.

   1.4 Recindx is defined as recname_INDX.

   Consider now the file which is parent to rec. Let it be denoted by fil.

   1.5 Set file_name to fil.

   1.6 If fil is of the form OLD.X or NEW.X set file_name

to OLD_X or NEW_X respectively and file_suff to file_nameU.

1.7 Otherwise set file_suff to file_nameS if the file is a source and to file_nameT if the file is a target.

1.8 Set eof to ENDFILE$file_name.

1.9 Retrieve the keyname associated with the record, if one exists, and assign it to key_name.

1.10 Set found to FOUND$file_name.

2. Issue the following declarations.

```
        DCL recbuf CHAR (len_dat(n)) ;

        DCL recindx FIXED BIN INIT(1) ;
```

This declares a buffer for the record into which and out of which the information will be read or written. 'Len_dat(n)' here gives the buffer length.

3. If the record is an output record, the instruction for moving the data from each field into the record buffer will be generated.

4. If the record is an output record and a SUBSET condition was specified for it we enclose the code for writing the record by the condition:

```
        IF SUBSET$rec THEN DO ;

                code

        END ;
```

The procedure DO_REC produces the code for reading and writing of records. It branches according to the cases in Table 7.1.

Table 7.1 The Various cases of program I/O control

Case 1: An Input Sequential and Nonkeyed Record.

The following code is produced:

```
IF $FSTfile_suff THEN DO ;
   READ FILE (file_suff) INTO (recbuf) ;
   $FSTfile_suff = '0'B ;
END ;
ELSE recbuf = filebuf ;
recindx = 1 ;
IF ^ENDFILE$file_name THEN
   READ FILE (file_suff) INTO (filebuf) ;
$ERROR_BUF = recbuf ;
```

The movement of the data to the individual fields will be done in conjunction with the nodes corresponding to the fields (see GENITEM). The next record is always read into file buffer so that we can unpack the data for the NEXT record.

Case 2: Input, Sequential and Keyed Record.

Ensure that the following reclarations have been issued:

```
DCL FOUND$rec BIT(1) ;
DCL PASSED$rec BIT(1) ;
```

Issue now the code:

```
FOUND$rec, PASSED$rec = '0'B ;
DO WHILE(^ENDFILE$file_name & ^PASSED$rec) ;
   READ FILE (file_suff) INTO (recbuf) ;
```

```
            (code for extracting the key field)
            IF keyname = POINTER$rec THEN
               FOUND$rec, PASSED$rec = '1'B ;
            ELSE IF keyname > POINTER$rec THEN
               PASSED$rec = '1'B ;
        END ;
        recindx = 1 ;
```

Case 3: Input, Nonsequential (ISAM), Keyed record.

Verify that the declaration

```
        DCL FOUND$rec BIT(1) ;
```

has been issued. Then issue the code:

```
        FOUND$rec = '1'B ;
        ON KEY (file_suff) FOUND$rec = '0'B ;
        READ FILE(file_suff) INTO(recbuf)
                KEY(POINTER$rec) ;
        recindx = 1 ;
```

Case 4: Output, Sequential Record.

Issue the following code:

```
        recindx = 1 ;
```

Call PACK procedure to pack its fields into the record
buffer. Then issue the code:

```
        WRITE FILE(file_suff) FROM(recbuf) ;
```

Case 5: Output, Nonsequential, Keyed and an Update Record
(both NEW and OLD specified)

Issue the following code:

```
        recindx = 1 ;
```

Call PACK procedure to pack its fields into the record buffer. Then issue the code:

```
        REWRITE FILE(file_suff) FROM(recbuf)
              KEY(POINTER$rec) ;
```

Case 6: Output, Nonsequential and Keyed Record.

.Issue the following code:

```
        recindx = 1 ;
```

Call PACK procedure to pack its fields into the record buffer. Then issue the code:

```
        WRITE FILE(file_suff) FROM(recbuf)
              KEY(POINTER$rec) ;
```

7.6   PACKING AND UNPACKING

After a record is read we unpack its fields from the record buffer and place them in the respective declared structures. Similarly before a record is written we pack its fields into the record buffer. The data movement is performed by individual transfers of fields. The transfer statements may be interleaved with other statements which control the iteration over respective fields' dimensions. The transfer instructions for unpacking are generated elsewhere, in conjunction with the schedule elements associated with the input field nodes. The code for packing an output record is generated in GENIOCD and inserted right

before the record buffer is to be written out.

7.6.1 PACK - PACKING THE OUTPUT FIELDS

The procedure PACK is called by GENIOCD in the case of an output record. It accepts a node number (NODE$) as input. It checks the type of the node NODE$. If the node is a field, it calls DO_FLD to generate the code for packing. Otherewise, it considers in turn each descendant of the node NODE$. For each descendant D it calls PACK1(D) recursively. PACK1: This procedure generates code for packing a node which may or may not repeat.

1. If the node is a repeating group or a field we get the termination criterion of the repetition.

    1.1 Open a loop: Call procedure GENDO to generate the DO-statement for opening the loop.

    1.2 Call the subprocedures PACK to issue code for packing a single element of the node.

    1.3 Call procedure GENEND to generate the code for terminating the loop.

2. If the node is not repeating then:
    Call procedure PACK to generate the code for packing all the constituent members of this node.

DO_FLD: This procedure is responsible for producing code to pack a field F into record buffer. It uses the procedure FIELDPK to generate the following code.

```
        SUBSTR(recbuf,recindx,lenstring) = F ;

        recindx = recindx+lenstring ;
```

FIELDPK is described further below.

7.6.2 GENITEM - UNPACKING THE INPUT FIELDS

This procedure is called to generate code for unpacking information from an input buffer to an input field. GEN_NODE calls GENITEM upon scanning a schedule element of an input field. GENITEM accepts as input the node number in the schedule element. The READ statement for reading the record to a buffer is generated by GENIOCD when the record node is scanned. GENITEM first finds for a record R the names of the input buffer RS and the packing counter RINDX. Next, GENITEM calls an auxiliary procedure FIELDPK, which generates the code for unpacking.

The GENITEM procedure is as follows:

1. Determine the name of the record containing the current field. Let it be rec. Then we construct a buffer name: rec_S and a buffer index name rec_INDX. Let the field's name be in the variable "field".

2. If the corresponding field in the next record is referenced, then call FIELDPK to unpack the field from the file buffer.

3. Call FIELDPK to generate the code for unpacking the field from the record buffer.

7.6.3 FIELDPK - PACKING AND UNPACKING FIELDS

The procedure FIELDPK produces the code for both the packing and unpacking operation. Input parameters are the field name, buffer name, record index name, and a code (CASE) to indicate whether the field has a NEXT prefix.

1. If the length type of the field is fixed, i.e. specified in the data description statements, we compute its length directly. If the field's type is 'C', 'N', or 'P', denoting respectively character, numeric or picture, we take the declared length. Otherwise we will compute the length of the field in bytes from its declared length and type. The string representing the length is stored in "lenstring".

2. If the length of the field was declared by specifying lower and upper bounds we check that there exists a control variable of the form LEN.field for this field. If none exists we issue the error message:

   FIELDPK: NO LENGTH SPECIFICATION FOR THE FIELD-field.

3. If a LEN.field control variable is found we set:

   lenstring = LEN.field

   The byte-length of the field will be computed during run time.

4. If the field is an input field we generate the instruction:

   UNSPEC(field) = SUBSTR(rec_S,rec_INDX,lenstring);

If the same field in the next record is referred in the specification, we will unpack the file buffer to get the corresponding field in the next record. For output field we generate:

    SUBSTR(rec_S,rec_INDX,lenstring) = UNSPEC(field) ;

Here "field" is the name properly subscripted and "lenstring" is the length specification. If the field is of type 'C', the UNSPEC qualifications will be omitted.

5. If the CASE code indicates that the field name does not have prefix NEXT then we generate the following code to update the buffer index:

    rec_INDX = rec_INDX+lenstring ;

There is no need to update recINDX if the unpacking is for a NEXT prefixed field.

7.7 GENERATING THE PROGRAM ERROR FILE

If a program error condition is induced during the execution of the generated program, then an input record, read during the iteration execution when the program error was induced is written to an error file, ERRORF. The required code for writing the bad input record to the error file is generated by the routines CODEGEN and GENIOCD. For example, the following PL/I code is included in PL1ON file:

    ON ERROR BEGIN ;

        WRITE FILE(ERRORF) FROM($ERROR_BUF) ;

```
                GO TO $ERR_LAB ;

            END ;
```

After the GENIOCD generate the code to read a record from an
input file it also generates a statement to copy the input
record into $ERROR_BUF.

7.8 GPL1DCL - GENERATING PL/I DECLARATION

This procedure generates the declarations for the data
nodes declared by the user and those added by the system.
As noted previously, some declarations are also generated by
other procedures during the code generation.

The main part of GPL1DCL is as follows:

1. For each file F in the specification (available from the
   list FILIST) call

        DECLARE_STRUCTURE(F)

   to declare F and all its descendants.

2. For each node N in the specification which is an interim
   variable or a control variable, call

        DECLARE_STRUCTURE(N)

3. For each subscript which has been used, issue the
   declaration:

        DCL subname FIXED BIN ;

7.8.1 DECLARE_STRUCTURE - DECLARING A STRUCTURE

This procedure is called by GPL1DCL. The input is a file node number. It declares the entire file structure. It issues the declarative: DECLARE, and then proceed to call DCL_STR(N,1,0).

7.8.1.1 DCL_STR(N, LEVEL, SUX)

This recursive procedure produces a declaring-clause for each node N in the structure. 'LEVEL' is the current level in the structure. SUX is a termination criterion stating whether there is a next node on the same level (younger brother) or a descendant.

1. Some Preliminary transformations are made on the declared node names.

    1.1 File names of the form NEW.F and OLD.F are modified to NEW_F and OLD_F respectively.

    1.2 The group names, record names, or field names are reduced to their stem (removing prefixes).

2. For control variables the resulting declaration is:

    For SIZE, and LEN names:

        name FIXED BIN,

    while for all other names:

        name BIT(1).

3. The declaration includes in general the following items:

        LEVEL - The component level.

Name   - The declared name.

Repetition - The number of physical storage elements.

Type   - The data type.

The data type is determined as follows:

For character fields - CHAR(len) [VARYING]

For numeric fields   - PIC '99....9'

For picture fields   - PIC 'picture'

For fixed binary   - BIN FIXED(len, scale)

For fixed decimal   - DEC FIXED(len,scale)

For binary floating   - BIN FLOAT(len)

For decimal floating - DEC FLOAT(len)

In the above 'len' is the specified or default length for the field. The VARYING option is taken if the length is specified (for strings) by a minimal length and a maximal length.

Repetition is defined in STOTYP of the node subscripts of the fields. If an array dimension is virtual we omit the repetition indicator. If an array dimension is a window of width k+1, the repetition is set to k+1. Otherwise, the array dimension must be a physical dimension. The node subscript list of the field node is scanned, and the repetition indicators for array dimensions are concatenated and put into a variable REP. If R is not an empty string, we will append the string '(REP)' after the declared field name.

4. For each of the descendants of the node M, call

DCL_STR(M,LEVEL+1,termination) recursively.

7.9 CGSUM - CODE GENERATION CONCLUSION

CGSUM has the task of concluding the code generation phase. First, the different files with the generated PL/I program (PL1DCL, PL1ON, PL1EX) are merged into one PL/I file (PL1PROG) which can be subsequently compiled. Secondly, a Code Generation Summary Report is written which lists the PL/I program. While the PL/I listing would not be of much use to the average MODEL user, it is of interest to the more sophisticated user and can serve the system programmer for insight or debugging of the MODEL system.

CHAPTER 8

SUGGESTED FUTURE RESEARCH

In this chapter we will discuss some of the possible directions of the future work. We have studied the issues related to analyzing the precedence relationships among the program events and ordering the program events to generate a program. There are additional techniques that need to be developed to reduce the execution time or the memory requirements. Two suggestions for program optimization area that require further research are described in this section.

8.1 ELIMINATING REDUNDANT COMPUTATION

8.1.1 ELIMINATING UNNECESSARY COPYING OF DATA

Consider the example of a stack which is represented by a pointer to the top of stack and a vector of elements. In defining a stack in the MODEL language it is necessary to define a new vector of elements each time when an element is

added to the top of the stack. Thus V(I,J) would be an array of vectors representing the stack and SIZE.V(I) would be the vector of pointers. The push function can be defined as

$$SIZE.V(I) = SIZE.V(I-1) + 1 ;$$

$$V(I,J) = IF\ J=SIZE.V(I)\ THEN\ new\text{-}element ;$$

$$ELSE\ V(I-1,J) ;$$

The copying in the ELSE part is very time-consuming when the stack is large. With our present program optimization approach, memory is allocated for two vectors V(I-1) and V(I), and the entire V(I-1,J) is copied into V(I,J). The suggested research would develop a method for recognizing the above illustrated condition and reducing both the memory required and execution time.

8.1.2 ELIMINATING MULTIPLE EVALUATIONS OF CONDITIONS

The assertions in the MODEL language may include conditions. In the case when the conditions in several statements are the same, it would be more efficient to form a block of the statements with the same condition and to execute the entire block only if the condition is true. A possible direction of future research is to recognize when condition expressions in several assertions are the same and to try during the scheduling to arrange these assertions in a block which will require only a single evaluation of the

condition.

In procedural languages the user can assemble statements within a BEGIN-END block and associate a condition expression with the entire block. In the MODEL system each, statement is scheduled by itself subject to a variety of considerations, including efficiency considerations. The suggestion here is to add an additional lower priority consideration whereby statements with the same condition expression will be placed in a block.

8.2 MODIFYING SPECIFICATION TO IMPROVE EFFICIENCY

A given computation task may be specified in a number of ways in the MODEL language. Since the program generated by the MODEL processor is influenced by the representation of the problem in a specification, different representations usually correspond to different programs. These programs may have different efficiency. For example consider the following MODEL specification. An input file IN contains a sequence of records, each with two fields called A(I) and B(I). The output is D, the quotient of dividing the sum of B's by the sum of A's. One way to state this problem in MODEL is to use P and C as interim variables as follows.

```
IN IS FILE (INREC(*)) ;
    INREC IS RECORD (A,B) ;
```

```
P = SUM(A(I),I) ;

C(I) = B(I)/P ;

D = SUM(C(I),I) ;
```

The generated program would scan the input file twice. In the first scan it computes the value of P and in the second scan value of D is computed. Since the input file is read only once in the generated program, we will have to save the whole file in main memory. However, there exists other MODEL specification which scan the input file only once and compute the same result. BY doing simple algebraic manipulation on the assertions, we can easily show that the following specification computes the same value of D.

```
IN IS FILE (INREC(*)) ;

  INREC IS RECORD (A,B) ;

P = SUM(A(I),I) ;

Q = SUM(B(I),I) ;

D = Q/P ;
```

This transformation on specification not only saves computation time but also the memory space. The goal of the transformation is to scan the input file only once so that there is no need to keep the whole file in the memory. If there is some computation which needs an input file and some other values which can be obtained after scanning the input file, then it is an indication that modifying the specification may be advantageous.

APPENDIX A

EXAMPLES OF MODEL SPECIFICATIONS

The appendix consists of two examples of MODEL specifications and the respective schedules generated by the system. These examples have been selected to illustrate the design decisions of the scheduler. The first example illustrates how the calculation of memory penalty effects the design of a schedule. The second example illustrates how the scope of an iteration may be enlarged based on analysis of related subscripts (i.e. through use of indirect subscripts.)

A.1 EXAMPLE OF TABLE LOOK-UP

This example consists of a bank customer file CUST which is updated based on a CODE which specifies the interest rate of each customer. The interest rates that correspond to codes are given in another input file TABLE. A new CUST file is produced with the updated balances. This is illustrated in Fig. A.1 and the MODEL specification is given in Fig. A.2.
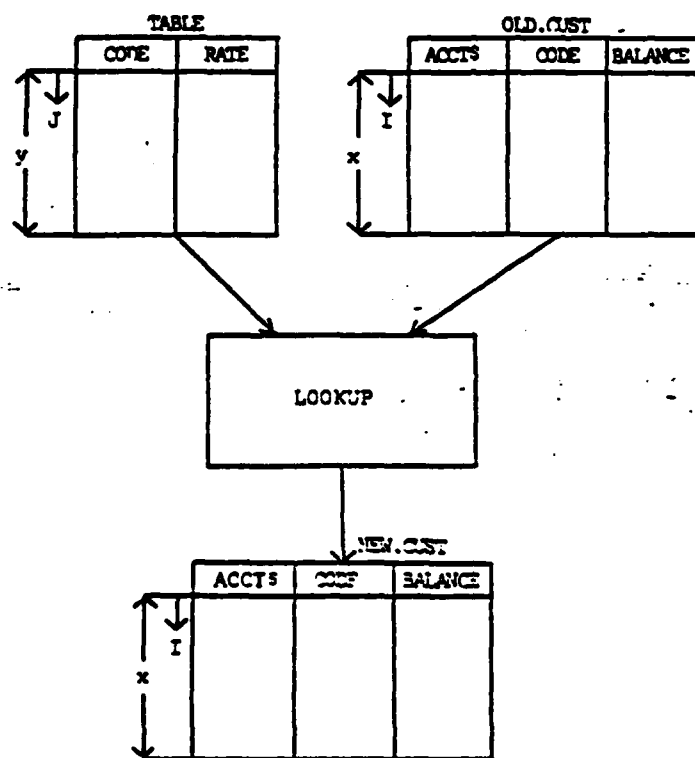

Fig. A.1 Diagram for the Example of LOOKUP

```
MODULE : LOOKUP;
SOURCE : CUST,TABLE;
TARGET : CUST;

CUST IS FILE (CUSTR(1:x));
    CUSTR IS RECORD (ACCT$,CODE,BALANCE);
        ACCT$ IS FIELD (NUMERIC(7));
        CODE IS FIELD (NUMERIC(4));
        BALANCE IS FIELD (PIC'(12)ZV.99');

TABLE IS FILE (TABLER(1:y));
    TABLER IS RECORD (CODE,RATE);
        CODE IS FIELD (NUMERIC(4));
        RATE IS FIELD (PIC'BV.99');

/********** ASSERSIONS FOR OUTPUT FILE CUSTOM **********/

NEW.ACCT$(I) = OLD.ACCT$(I);

NEW.CODE(I) = OLD.CODE(I);

IF TABLE.CODE(J)=OLD.CODE(I) THEN
    NEW.BALANCE(I) =
              OLD.BALANCE(I) * (1 + RATE(J));

END.TABLER=ENDFILE.TABLER;

END.OLD.CUSTR=ENDFILE.OLD.CUSTR;

/********** END OF THE SPECIFICATION ***************/
```

Fig. A.2 MODEL specification for LOOKUP

The most efficient memory usage depends on the relative sizes of TABLE and CUST, i.e. on x and y respectively. Only one of these files can have a virtual memory allocation. If TABLE is relatively very large, then it should have virtual memory allocation and CUST must then have a physical memory allocation, and vice versa if CUST is the larger file.

Fig. A.3 shows the Array Graph with the two alternative
range sets that are candidates for a loop scope circled.
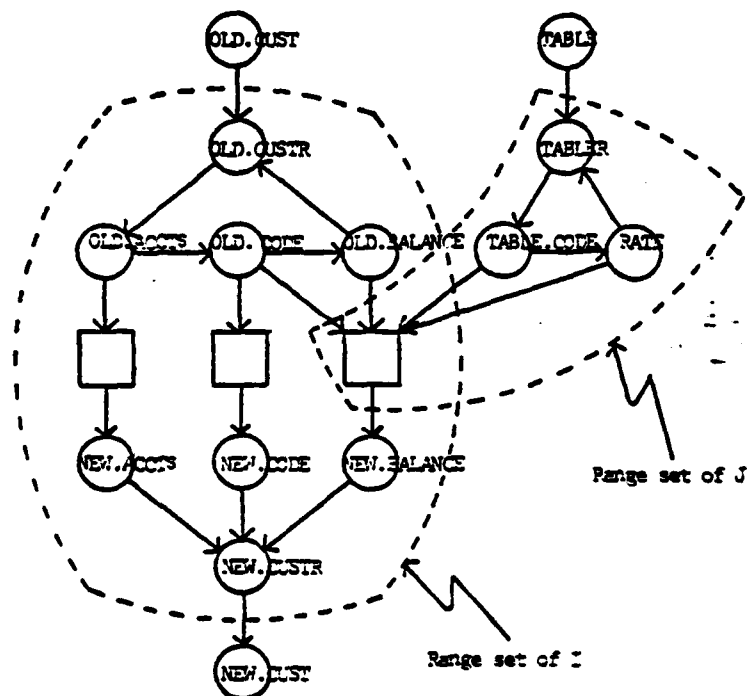The memory penalties for these two alternatives are as
follows.

Fig. A.3 Array Graph of the LOOKUP specification

If a loop iterated over the first range, i.e. I, is
scheduled first, then three arrays have to become physical,
i.e. END.TABLER, RATE, and TABLE.CODE. END.OLD.CUSTR has
to be a window of width two. The memory penalty is computed
as follows:

```
END.TABLER      : (y - 1) * 1 = y-1

RATE            : (y - 1) * 4 = 4y-4

TABLE.CODE      : (y - 1) * 4 = 4y-4

END.OLD.CUSTR   : (2 - 1) * 1 = 1
```

total penalty = 9y-8

If a loop iterated over the second range, i.e.   J, is scheduled first, then four arrays have to become physical, i.e. OLD.CODE, OLD.BALANCE, NEW.BALANCE, and END.OLD.CUSTR. END.TABLER has to be a window of width two. The memory penalty is computed as follows:

```
OLD.CODE        : (x - 1) * 4  = 4x-4

OLD.BALANCE     : (x - 1) * 15 = 15x-15

NEW.BALANCE     : (x - 1) * 15 = 15x-15

END.OLD.CUSTR   : (x - 1) * 1  = x-1

END.TABLER      : (2 - 1) * 1  = 1
```

total penalty = 35x-34

Depending on the relative values of x and y the scheduler may produce the schedules in Fig. A.4 and Fig. A.6. In the case that x is equal to 10 and y equal to 37, the TABLE file is relatively larger, the system will make it virtual. If TABLE is the virtual (larger) file, then the schedule has first an iteration for reading in OLD.CUST. Next an iteration reads one record of TABLE at a

time and computes NEW.BALANCE ᵣ r all customers with the respective CODE. Finally a third iteration writes out the NEW.CUST file. The corresponding PL/I program generated by the system is listed in Fig. A.5.

In the case that x is equal to 10 and y equal to 35, the CUST file is relatively larger, the system will make the CUST file virtual. The schedule in Fig. A.6, when CUST is virtual, has first an iteration for reading TABLE. This is followed by an iteration for reading, updating, and writing a record of CUST at a time. The corresponding PL/I program generated by the system is listed in Fig. A.7.
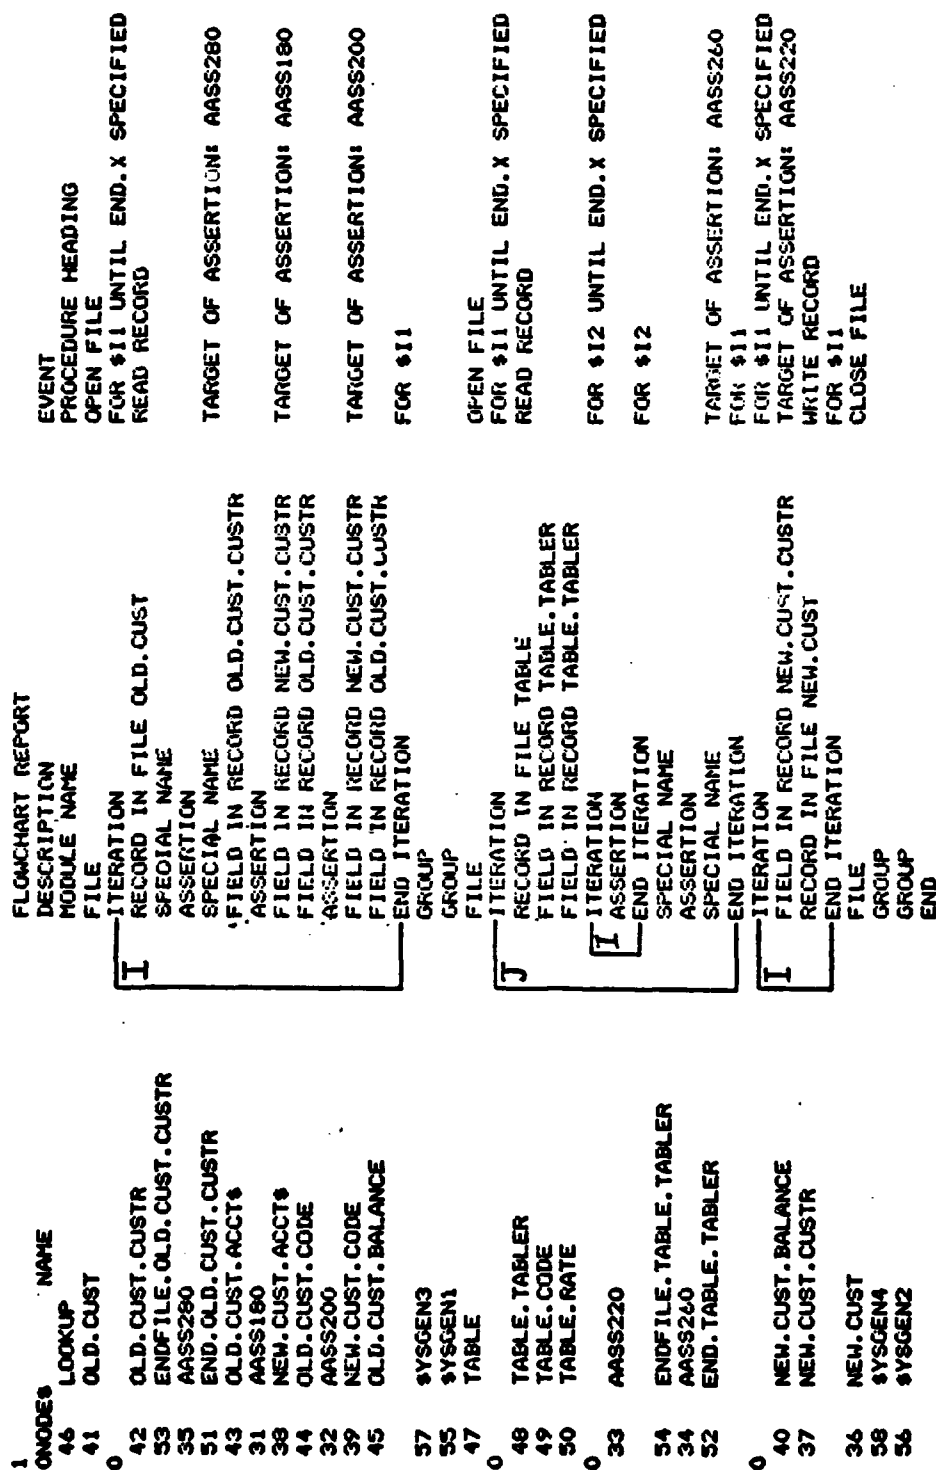
FLOWCHART REPORT

ONODES	NAME	DESCRIPTION	EVENT
1			
46	LOOKUP	MODULE NAME	PROCEDURE HEADING
41	OLD.CUST	FILE	OPEN FILE
0		ITERATION	FOR $I1 UNTIL END.X SPECIFIED
42	OLD.CUST.CUSTR	RECORD IN FILE OLD.CUST	READ RECORD
53	ENDFILE.OLD.CUST.CUSTR	SPECIAL NAME	
35	AASS280	ASSERTION	
51	END.OLD.CUST.CUSTR	SPECIAL NAME	TARGET OF ASSERTION: AASS280
43	OLD.CUST.ACCT$	FIELD IN RECORD OLD.CUST.CUSTR	
31	AASS180	ASSERTION	
38	NEW.CUST.ACCT$	FIELD IN RECORD NEW.CUST.CUSTR	TARGET OF ASSERTION: AASS180
44	OLD.CUST.CODE	FIELD IN RECORD OLD.CUST.CUSTR	
32	AASS200	ASSERTION	
39	NEW.CUST.CODE	FIELD IN RECORD NEW.CUST.CUSTR	TARGET OF ASSERTION: AASS200
45	OLD.CUST.BALANCE	FIELD IN RECORD OLD.CUST.CUSTR	
		END ITERATION	FOR $I1
57	SYSGEN3	GROUP	
55	SYSGEN1	GROUP	
47	TABLE	FILE	OPEN FILE
0		ITERATION	FOR $I1 UNTIL END.X SPECIFIED
48	TABLE.TABLER	RECORD IN FILE TABLE	READ RECORD
49	TABLE.CODE	FIELD IN RECORD TABLE.TABLER	
50	TABLE.RATE	FIELD IN RECORD TABLE.TABLER	
		ITERATION	FOR $I2 UNTIL END.X SPECIFIED
33	AASS220	ASSERTION	
		END ITERATION	FOR $I2
54	ENDFILE.TABLE.TABLER	SPECIAL NAME	
34	AASS260	ASSERTION	
52	END.TABLE.TABLER	SPECIAL NAME	TARGET OF ASSERTION: AASS260
		END ITERATION	FOR $I1
0		ITERATION	FOR $I1 UNTIL END.X SPECIFIED
40	NEW.CUST.BALANCE	FIELD IN RECORD NEW.CUST.CUSTR	TARGET OF ASSERTION: AASS220
37	NEW.CUST.CUSTR	RECORD IN FILE NEW.CUST	WRITE RECORD
		END ITERATION	FOR $I1
36	NEW.CUST	FILE	CLOSE FILE
58	SYSGEN4	GROUP	
56	SYSGEN2	GROUP	
		END	

Fig. A.4 Schedule-1 for the Example of LOOKUP

```
LOOKUP: PROCEDURE OPTIONS(MAIN);
DCL CUSTS RECORD SEQL INPUT;
DCL $FSTCUSTS BIT(1) INIT('1'B);
DCL ENDFILE$CUSTS BIT(1) INIT('O'B);
DCL OLD_CUST_S CHAR(26) VARYING INIT('');
DCL OLD_CUST_INDX FIXED BIN;
DCL OLD_CUST_CUSTR_S CHAR(26) VARYING;
OCL OLD_CUST_CUSTR_INDX FIXED BIN;
DCL TABLES RECORD SEQL INPUT;
DCL $FSTTABLES BIT(1) INIT('1'B);
DCL ENDFILE$TABLES BIT(1) INIT('O'B);
OCL TABLE_S CHAR(3) VARYING INIT('');
DCL TABLE_INDX FIXED BIN;
DCL TABLE_TABLER_S CHAR(3) VARYING;
DCL TABLE_TABLER_INDX FIXED BIN;
DCL NEW_CUST_CUSTR_S CHAR(26) VARYING;
DCL NEW_CUST_CUSTR_S_F CHAR(26);
DCL NEW_CUST_CUSTR_SC BIT(208) BASED(ADDR(NEW_CUST_CUSTR_S_F));
DCL NEW_CUST_CUSTR_INDX FIXED BIN;
DCL CUSTT RECORD SEQL OUTPUT;
DCL $FSTCUSTT BIT(1) INIT('1'B);
DCL $ERROR_BUF CHAR(270) VAR;
DCL ERRORF FILE RECORD OUTPUT;
OCL ERRORF_BIT BIT(1) STATIC INIT('1'B);
DCL ($ERROR,$ACC_ERROR,$NOT_DONE)(20) BIT(1);
OCL $ERR_LAB(20) LABEL;
DCL $ERRSP% FIXED BIN STATIC INITIAL (0);
OCL $TMP_VAL FLOAT BIN;
DCL $TMP_ERR BIT(1);
DECLARE
    1 NEW_CUST,
     2 CUSTR,
      3 ACCTS(10) PIC'9999999',
      3 CODE(10) PIC'9999',
      3 BALANCE(10) PIC'(12)ZV.99';
DECLARE
    1 OLD_CUST,
     2 CUSTR,
      3 ACCTS PIC'9999999',
      3 CODE(10) PIC'9999',
      3 BALANCE(10) PIC'(12)ZV.99';
DECLARE
    1 TABLE,
     2 TABLER,
      3 CODE PIC'9999',
      3 RATE PIC'9V.99';
DECLARE
    1 INTERIM,
     2 $YSGEN1,
      3 END$OLD_CUST_CUSTR(10) BIT(1) ,
     2 $YSGEN2,
      3 END$TABLE_TABLER(2) BIT(1) ,
     2 $YSGEN3,
      3 ENDFILE$OLD_CUST_CUSTR BIT(1) ,
     2 $YSGEN4,
      3 ENDFILE$TABLE_TABLER BIT(1) ;
DCL $I1 FIXED BIN;
DCL $I2 FIXED BIN;
DCL (TRUE,SELECTED) BIT(1) INIT('1'B);
DCL (FALSE,NOT_SELE,NOT_SELECTED) BIT(1) INIT('O'B);
```

Fig. A.5 Generated PL/I program for Schedule-1

```
ON ENDFILE(CUSTS) BEGIN;
ENDFILE$CUSTS='1'B;
 OLD_CUST_S=COPY(' ',26);
 END;
ON ENDFILE(TABLES) BEGIN;
ENDFILE$TABLES='1'B;
 TABLE_S=COPY(' ',8);
 END;
ON UNDEFINEDFILE(ERRORF)  ERRORF_BIT='0'B;
/* ON ERROR BEGIN;
IF ERRORF_BIT THEN WRITE FILE(ERRORF) FROM ($ERROR_BUF);
$ERROR($ERRSP$)='1'B;
$ACC_ERROR($ERRSP$)='1'B;
GO TO $ERR_LAB($ERRSP$) ;
 END ;
*/ ;
ERROR_RESTART: ;
$ERRSP$ = $ERRSP$ +1;
$ERROR($ERRSP$)='0'B;
$ACC_ERROR($ERRSP$)='0'B;
$ERR_LAB($ERRSP$)=END_PROGRAM;
OPEN FILE(CUSTS);
$ERRSP$ = $ERRSP$ +1;
$ERROR($ERRSP$)='0'B;
$ACC_ERROR($ERRSP$)='0'B;
$ERR_LAB($ERRSP$)=LOOP_END1;
$I1 =0;
$NOT_DONE(1)='1'B;
DO  WHILE($NOT_DONE(1));
  $I1 = $I1 +1;
  $ERROR($ERRSP$)='0'B;
  IF $FSTCUSTS THEN DO;
   READ FILE(CUSTS) INTO (OLD_CUST_CUSTR_S);
   $FSTCUSTS='0'B;
   END;
  ELSE OLD_CUST_CUSTR_S=OLD_CUST_S;
  OLD_CUST_CUSTR_INDX=1;
  IF ^ENDFILE$CUSTS THEN READ FILE(CUSTS) INTO (OLD_CUST_S);
  $ERROR_BUF=OLD_CUST_CUSTR_S;
  ENDFILE$OLD_CUST_CUSTR=ENDFILE$CUSTS;
  END$OLD_CUST_CUSTR($I1)=ENDFILE$OLD_CUST_CUSTR;
  UNSPEC(OLD_CUST.ACCT$)=UNSPEC(SUBSTR(OLD_CUST_CUSTR_S,OLD_CUST_CUSTR_INDX,7))
  ;
  OLD_CUST_CUSTR_INDX=OLD_CUST_CUSTR_INDX+7 ;
  NEW_CUST.ACCT$($I1)=OLD_CUST.ACCT$;
  UNSPEC(OLD_CUST.CODE($I1))=UNSPEC(SUBSTR(OLD_CUST_CUSTR_S,OLD_CUST_CUSTR_INDX,
  4)) ;
  OLD_CUST_CUSTR_INDX=OLD_CUST_CUSTR_INDX+4 ;
  NEW_CUST.CODE($I1)=OLD_CUST.CODE($I1);
  UNSPEC(OLD_CUST.BALANCE($I1))=UNSPEC(SUBSTR(OLD_CUST_CUSTR_S,
  OLD_CUST_CUSTR_INDX,15)) ;
  OLD_CUST_CUSTR_INDX=OLD_CUST_CUSTR_INDX+15 ;
  LOOP_END1: ;
  IF END$OLD_CUST_CUSTR($I1) THEN $NOT_DONE(1)='0'B;
END;
$TMP_ERR=$ACC_ERROR($ERRSP$);
$ERRSP$ = $ERRSP$ - 1;
IF $TMP_ERR THEN $ERROR($ERRSP$)='1'B;
IF $TMP_ERR THEN $ACC_ERROR($ERRSP$)='1'B;
OPEN FILE(TABLES);
$ERRSP$ = $ERRSP$ +1;
```

Fig. A.5 Generated PL/I program for Schedule-1

(Continued)

```
$ERROR($ERRSP$)='0'B$
$ACC_ERROR($ERRSP$)='0'B$
$ERR_LAB($ERRSP$)=LOOP_END2$
$I1 =0$
$NOT_DONE(1)='1'B$
DO  WHILE($NOT_DONE(1))$
  $I1 = $I1 +1$
  $ERROR($ERRSP$)='0'B$
  IF $FSTTABLES THEN DO$
   READ FILE(TABLES) INTO (TABLE_TABLER_S)$
   $FSTTABLES='0'B$
   END$
  ELSE TABLE_TABLER_S=TABLE_S$
  TABLE_TABLER_INDX=1$
  IF ^ENDFILE$TABLES THEN READ FILE(TABLES) INTO (TABLE_S)$
  $ERROR_BUF=TABLE_TABLER_S$
  ENDFILE$TABLE_TABLER=ENDFILE$TABLES$
  UNSPEC(TABLE.CODE)=UNSPEC(SUBSTR(TABLE_TABLER_S,TABLE_TABLER_INDX,4)) $
  TABLE_TABLER_INDX=TABLE_TABLER_INDX+4 $
  UNSPEC(TABLE.RATE)=UNSPEC(SUBSTR(TABLE_TABLER_S,TABLE_TABLER_INDX,4)) $
  TABLE_TABLER_INDX=TABLE_TABLER_INDX+4 $
  $I2 =0$
  $NOT_DONE(2)='1'B$
  DO  WHILE($NOT_DONE(2))$
    $I2 = $I2 +1$
    IF TABLE.CODE=OLD_CUST.CODE($I2) THEN NEW_CUST.BALANCE($I2)=
    OLD_CUST.BALANCE($I2)*(1+TABLE.RATE)$
    LOOP_END3$$
    IF END$OLD_CUST_CUSTR($I2) THEN $NOT_DONE(2)='0'B$
  END$
  END$TABLE_TABLER(2)=ENDFILE$TABLE_TABLER$
  LOOP_END2$$
  IF END$TABLE_TABLER(2) THEN $NOT_DONE(1)='0'B$
  END$TABLE_TABLER(1) = END$TABLE_TABLER(2)$
END$
$TMP_ERR=$ACC_ERROR($ERRSP$)$
$ERRSP$ = $ERRSP$ - 1$
IF $TMP_ERR THEN $ERROR($ERRSP$)='1'B$
IF $TMP_ERR THEN $ACC_ERROR($ERRSP$)='1'B$
$I1 =0$
$NOT_DONE(1)='1'B$
DO  WHILE($NOT_DONE(1))$
  $I1 = $I1 +1$
  NEW_CUST_CUSTR_INDX=1$
  SUBSTR(NEW_CUST_CUSTR_SC,NEW_CUST_CUSTR_INDX*8-7,7*8)=UNSPEC(NEW_CUST.ACCT$(
  $I1)) $
  NEW_CUST_CUSTR_INDX=NEW_CUST_CUSTR_INDX+7 $
  SUBSTR(NEW_CUST_CUSTR_SC,NEW_CUST_CUSTR_INDX*8-7,4*8)=UNSPEC(NEW_CUST.CODE(
  $I1)) $
  NEW_CUST_CUSTR_INDX=NEW_CUST_CUSTR_INDX+4 $
  SUBSTR(NEW_CUST_CUSTR_SC,NEW_CUST_CUSTR_INDX*8-7,15*8)=UNSPEC(
  NEW_CUST.BALANCE($I1)) $
  NEW_CUST_CUSTR_INDX=NEW_CUST_CUSTR_INDX+15 $
  NEW_CUST_CUSTR_S=SUBSTR(NEW_CUST_CUSTR_S_F,1,NEW_CUST_CUSTR_INDX-1)$
  WRITE FILE(CUSTT) FROM (NEW_CUST_CUSTR_S)$
  LOOP_END4$$
  IF END$OLD_CUST_CUSTR($I1) THEN $NOT_DONE(1)='0'B$
END$
CLOSE FILE(CUSTT)$
END_PROGRAM: RETURN$
END LOOKUP$
$
```

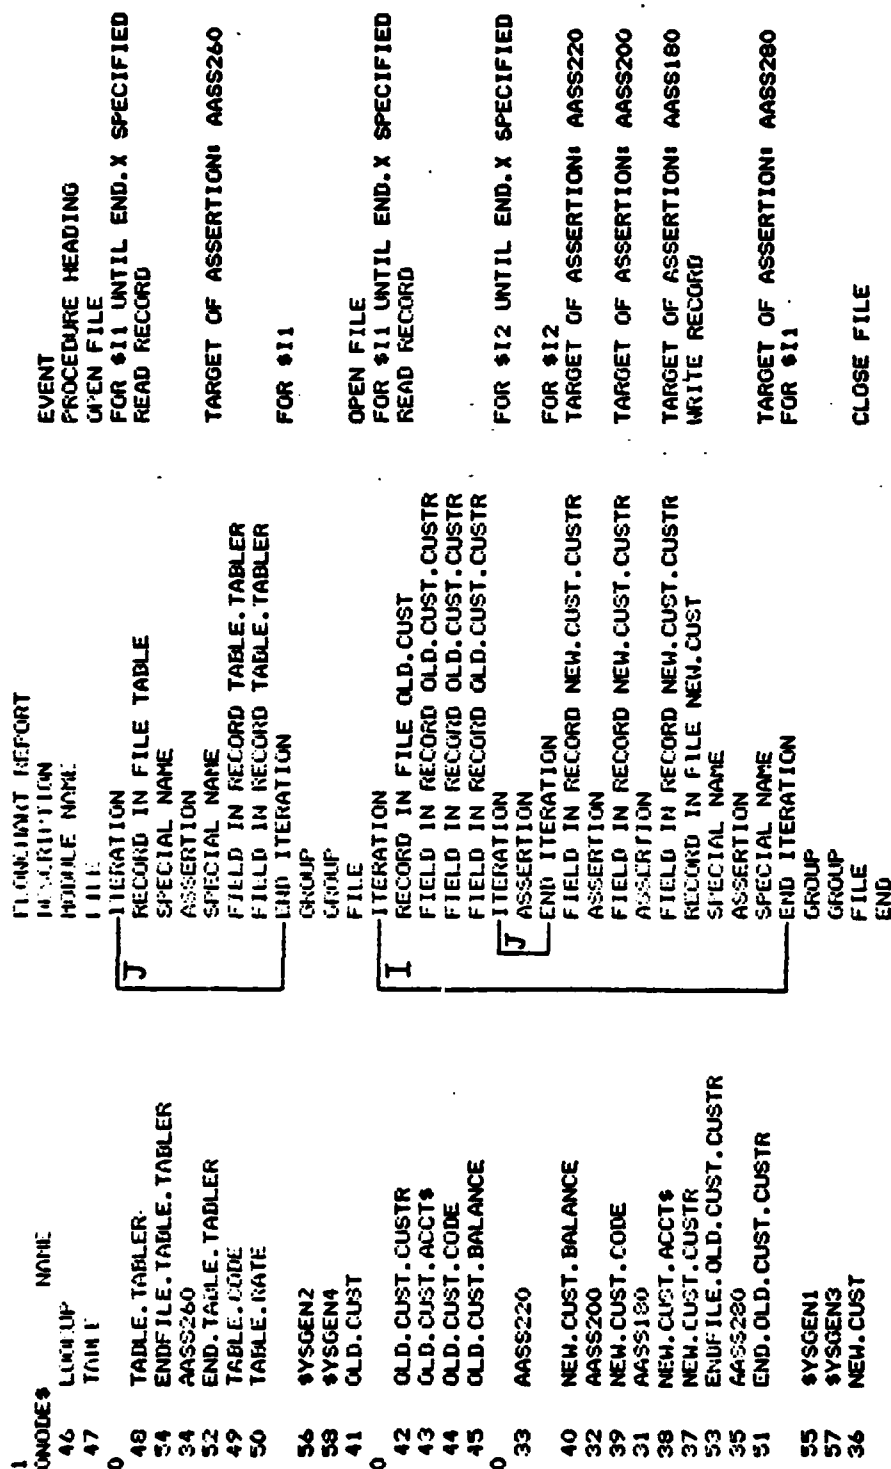Fig. A.5 Generated PL/I program for Schedule-1
(Continued)

ONODE# NAME

46 LOOKUP
47 TIME

48 TABLE.TABLER.
54 ENDFILE.TABLE.TABLER
34 AASS260
52 END.TABLE.TABLER
49 TABLE.CODE
50 TABLE.RATE

56 $YSGEN2
58 $YSGEN4
41 OLD.CUST

42 OLD.CUST.CUSTR
43 OLD.CUST.ACCT$
44 OLD.CUST.CODE
45 OLD.CUST.BALANCE

33 AASS220

40 NEW.CUST.BALANCE
32 AASS200
39 NEW.CUST.CODE
31 AASS180
38 NEW.CUST.ACCT$
37 NEW.CUST.CUSTR
53 ENDFILE.OLD.CUST.CUSTR
35 AASS280
51 END.OLD.CUST.CUSTR

55 $YSGEN1
57 $YSGEN3
36 NEW.CUST

FLOWCHART REPORT
DESCRIPTION
MODULE NAME
TITLE
ITERATION
RECORD IN FILE TABLE
SPECIAL NAME
ASSERTION
SPECIAL NAME
FIELD IN RECORD TABLE.TABLER
FIELD IN RECORD TABLE.TABLER
END ITERATION
GROUP
GROUP
FILE
ITERATION
RECORD IN FILE OLD.CUST
FIELD IN RECORD OLD.CUST.CUSTR
FIELD IN RECORD OLD.CUST.CUSTR
FIELD IN RECORD OLD.CUST.CUSTR
ITERATION
ASSERTION
END ITERATION
FIELD IN RECORD NEW.CUST.CUSTR
ASSERTION
FIELD IN RECORD NEW.CUST.CUSTR
ASSERTION
FIELD IN RECORD NEW.CUST.CUSTR
RECORD IN FILE NEW.CUST
SPECIAL NAME
ASSERTION
SPECIAL NAME
END ITERATION
GROUP
GROUP
FILE
END

EVENT
PROCEDURE HEADING
OPEN FILE
FOR $I1 UNTIL END.X SPECIFIED
READ RECORD

TARGET OF ASSERTION: AASS260

FOR $I1

OPEN FILE
FOR $I1 UNTIL END.X SPECIFIED
READ RECORD

FOR $I2 UNTIL END.X SPECIFIED

FOR $I2
TARGET OF ASSERTION: AASS220

TARGET OF ASSERTION: AASS200

TARGET OF ASSERTION: AASS180
WRITE RECORD

TARGET OF ASSERTION: AASS280
FOR $I1

CLOSE FILE

Fig. A.6 Schedule-2 for the Example of LOOKUP

```
LOOKUP: PROCEDURE OPTIONS(MAIN);
DCL TABLES RECORD SEQL INPUT;
DCL $FSTTABLES BIT(1) INIT('1'B);
DCL ENDFILE$TABLES BIT(1) INIT('0'B);
DCL TABLE_S CHAR(8) VARYING INIT('');
DCL TABLE_INDX FIXED BIN;
DCL TABLE_TABLER_S CHAR(8) VARYING;
DCL TABLE_TABLER_INDX FIXED BIN;
DCL CUSTS RECORD SEQL INPUT;
DCL $FSTCUSTS BIT(1) INIT('1'B);
DCL ENDFILE$CUSTS BIT(1) INIT('0'B);
DCL OLD_CUST_S CHAR(26) VARYING INIT('');
DCL OLD_CUST_INDX FIXED BIN;
DCL OLD_CUST_CUSTR_S CHAR(26) VARYING;
DCL OLD_CUST_CUSTR_INDX FIXED BIN;
DCL NEW_CUST_CUSTR_S CHAR(26) VARYING;
DCL NEW_CUST_CUSTR_S_F CHAR(26);
DCL NEW_CUST_CUSTR_SC BIT(208) BASED(ADDR(NEW_CUST_CUSTR_S_F));
DCL NEW_CUST_CUSTR_INDX FIXED BIN;
DCL CUSTT RECORD SEQL OUTPUT;
DCL $FSTCUSTT BIT(1) INIT('1'B);
DCL $ERROR_BUF CHAR(270) VAR;
DCL ERRORF FILE RECORD OUTPUT;
DCL ERRORF_BIT BIT(1) STATIC INIT('1'B);
DCL ($ERROR,$ACC_ERROR,$NOT_DONE)(20) BIT(1);
DCL $ERR_LAB(20) LABEL;
DCL $ERRSP$ FIXED BIN STATIC INITIAL (0);
DCL $TMP_VAL FLOAT BIN;
DCL $TMP_ERR BIT(1);
DECLARE
    1 NEW_CUST,
    2 CUSTR,
      3 ACCTS PIC'9999999',
      3 CODE PIC'9999',
      3 BALANCE PIC'(12)ZV.99';
DECLARE
    1 OLD_CUST,
    2 CUSTR,
      3 ACCTS PIC'9999999',
      3 CODE PIC'9999',
      3 BALANCE PIC'(12)ZV.99';
DECLARE
    1 TABLE,
    2 TABLER,
      3 CODE(35) PIC'9999',
      3 RATE(35) PIC'9V.99';
DECLARE
    1 INTERIM,
    2 SYSGEN1,
      3 END$OLD_CUST_CUSTR(2) BIT(1) ,
    2 SYSGEN2,
      3 END$TABLE_TABLER(35) BIT(1) ,
    2 SYSGEN3,
      3 ENDFILE$OLD_CUST_CUSTR BIT(1) ,
    2 SYSGEN4,
      3 ENDFILE$TABLE_TABLER BIT(1) ;
DCL $I1 FIXED BIN;
DCL $I2 FIXED BIN;
DCL (TRUE,SELECTED) BIT(1) INIT('1'B);
DCL (FALSE,NOT_SELE,NOT_SELECTED) BIT(1) INIT('0'B);
```

Fig. A.7 Generated PL/I program for Schedule-2

```
ON ENDFILE(TABLES) BEGIN;
ENDFILE$TABLES='1'B;
 TABLE_S=COPY(' ',8);
 END;
ON ENDFILE(CUSTS) BEGIN;
ENDFILE$CUSTS='1'B;
 OLD_CUST_S=COPY(' ',26);
 END;
ON UNDEFINEDFILE(ERRORF)  ERRORF_BIT='0'B;
/* ON ERROR BEGIN;
IF ERRORF_BIT THEN WRITE FILE(ERRORF) FROM ($ERROR_BUF);
$ERROR($ERRSP$)='1'B;
$ACC_ERROR($ERRSP$)='1'B;
GO TO $ERR_LAB($ERRSP$) ;
 END ;
*/;
ERROR_RESTART:  ;
$ERRSP$ = $ERRSP$ +1;
$ERROR($ERRSP$)='0'B;
$ACC_ERROR($ERRSP$)='0'B;
$ERR_LAB($ERRSP$)=END_PROGRAM;
OPEN FILE(TABLES);
$ERRSP$ = $ERRSP$ +1;
$ERROR($ERRSP$)='0'B;
$ACC_ERROR($ERRSP$)='0'B;
$ERR_LAB($ERRSP$)=LOOP_END1;
$I1 =0;
$NOT_DONE(1)='1'B;
DO  WHILE($NOT_DONE(1));
  $I1 = $I1 +1;
   $ERROR($ERRSP$)='0'B;
   IF $FST$TABLES THEN DO;
    READ FILE(TABLES) INTO (TABLE_TABLER_S);
    $FST$TABLES='0'B;
    END;
   ELSE TABLE_TABLER_S=TABLE_S;
   TABLE_TABLER_INDX=1;
   IF `ENDFILE$TABLES THEN READ FILE(TABLES) INTO (TABLE_S);
   $ERROR_BUF=TABLE_TABLER_S;
   ENDFILE$TABLE_TABLER=ENDFILE$TABLES;
   END$TABLE_TABLER($I1)=ENDFILE$TABLE_TABLER;
   UNSPEC(TABLE.CODE($I1))=UNSPEC(SUBSTR(TABLE_TABLER_S,TABLE_TABLER_INDX,4)) ;
   TABLE_TABLER_INDX=TABLE_TABLER_INDX+4 ;
   UNSPEC(TABLE.RATE($I1))=UNSPEC(SUBSTR(TABLE_TABLER_S,TABLE_TABLER_INDX,4)) ;
   TABLE_TABLER_INDX=TABLE_TABLER_INDX+4 ;
   LOOP_END1:;
   IF END$TABLE_TABLER($I1) THEN $NOT_DONE(1)='0'B;
END;
$TMP_ERR=$ACC_ERROR($ERRSP$);
$ERRSP$ = $ERRSP$ - 1;
IF $TMP_ERR THEN $ERROR($ERRSP$)='1'B;
IF $TMP_ERR THEN $ACC_ERROR($ERRSP$)='1'B;
OPEN FILE(CUSTS);
$ERRSP$ = $ERRSP$ +1;
$ERROR($ERRSP$)='0'B;
$ACC_ERROR($ERRSP$)='0'B;
$ERR_LAB($ERRSP$)=LOOP_END2;
$I1 =0;
$NOT_DONE(1)='1'B;
```

Fig. A.7 Generated PL/I program for Schedule-2
(Continued)

```
DO  WHILE($NOT_DONE(1));
  $I1 = $I1 +1;
  $ERROR($ERRSP$)='0'B;
  IF $FSTCUSTS THEN DO;
   READ FILE(CUSTS) INTO (OLD_CUST_CUSTR_S);
   $FSTCUSTS='0'B;
   END;
  ELSE OLD_CUST_CUSTR_S=OLD_CUST_S;
  OLD_CUST_CUSTR_INDX=1;
  IF ^ENDFILE$CUSTS THEN READ FILE(CUSTS) INTO (OLD_CUST_S);
  $ERROR_BUF=OLD_CUST_CUSTR_S;
  ENDFILE$OLD_CUST_CUSTR=ENDFILE$CUSTS;
  UNSPEC(OLD_CUST.ACCT$)=UNSPEC(SUBSTR(OLD_CUST_CUSTR_S,OLD_CUST_CUSTR_INDX,7))
  ;
  OLD_CUST_CUSTR_INDX=OLD_CUST_CUSTR_INDX+7 ;
  UNSPEC(OLD_CUST.CODE)=UNSPEC(SUBSTR(OLD_CUST_CUSTR_S,OLD_CUST_CUSTR_INDX,4)) ;
  OLD_CUST_CUSTR_INDX=OLD_CUST_CUSTR_INDX+4 ;
  UNSPEC(OLD_CUST.BALANCE)=UNSPEC(SUBSTR(OLD_CUST_CUSTR_S,OLD_CUST_CUSTR_INDX,
  15)) ;
  OLD_CUST_CUSTR_INDX=OLD_CUST_CUSTR_INDX+15 ;
  $I2 =0;
  $NOT_DONE(2)='1'B;
  DO  WHILE($NOT_DONE(2));
    $I2 = $I2 +1;
    IF TABLE.CODE($I2)=OLD_CUST.CODE THEN NEW_CUST.BALANCE=OLD_CUST.BALANCE*(
    1+TABLE.RATE($I2));
    LOOP_END3; ;
    IF END$TABLE_TABLER($I2) THEN $NOT_DONE(2)='0'B;
  END;
  NEW_CUST.CODE=OLD_CUST.CODE;
  NEW_CUST.ACCT$=OLD_CUST.ACCT$;
  NEW_CUST_CUSTR_INDX=1;
  SUBSTR(NEW_CUST_CUSTR_SC,NEW_CUST_CUSTR_INDX*8-7,7*8)=UNSPEC(NEW_CUST.ACCT$) ;
  NEW_CUST_CUSTR_INDX=NEW_CUST_CUSTR_INDX+7 ;
  SUBSTR(NEW_CUST_CUSTR_SC,NEW_CUST_CUSTR_INDX*8-7,4*8)=UNSPEC(NEW_CUST.CODE) ;
  NEW_CUST_CUSTR_INDX=NEW_CUST_CUSTR_INDX+4 ;
  SUBSTR(NEW_CUST_CUSTR_SC,NEW_CUST_CUSTR_INDX*8-7,15*8)=UNSPEC(
  NEW_CUST.BALANCE) ;
  NEW_CUST_CUSTR_INDX=NEW_CUST_CUSTR_INDX+15 ;
  NEW_CUST_CUSTR_S=SUBSTR(NEW_CUST_CUSTR_S_F,1,NEW_CUST_CUSTR_INDX-1);
  WRITE FILE(CUSTT) FROM (NEW_CUST_CUSTR_S);
  END$OLD_CUST_CUSTR(2)=ENDFILE$OLD_CUST_CUSTR;
  LOOP_END2; ;
  IF END$OLD_CUST_CUSTR(2) THEN $NOT_DONE(1)='0'B;
  END$OLD_CUST_CUSTR(1) = END$OLD_CUST_CUSTR(2);
END;
$TMP_ERR=$ACC_ERROR($ERRSP$);
$ERRSP$ = $ERRSP$ - 1;
IF $TMP_ERR THEN $ERROR($ERRSP$)='1'B;
IF $TMP_ERR THEN $ACC_ERROR($ERRSP$)='1'B;
CLOSE FILE(CUSTT);
END_PROGRAM: RETURN;
END LOOKUP;
```

Fig. A.7 Generated PL/I program for Schedule-2
(Continued)

A.2 EXAMPLE OF MERGE OF FOUR FILES

This example illustrates merging the scopes of loops
for related subscripts, thus increasing the scope of loops,
decreasing the number of loops in a program, and permitting
virtual memory allocation for arrays referenced in the
merged loops. The example shows also how this merging can
be applied recursively, increasing the scope of loops on
every application. It consists of merger of four files,
first merging two pairs, S1 and S2 into M1, and S3 and S4
into M2, and then merging M1 and M2 into T.

This is illustrated in Fig. A.8. Each of the files
consists of records R, each with two fields, NUM and CHR.
The records in each files are sorted by increased values of
NUM. The three merger boxes in Fig. A.8 are similar and it
suffices to show only the merger of S1 and S2 into M1. The
respective specification and Array Graph are shown in
Fig. A.9 and A.10. The range sets in Fig. A.10 are shown
circled.
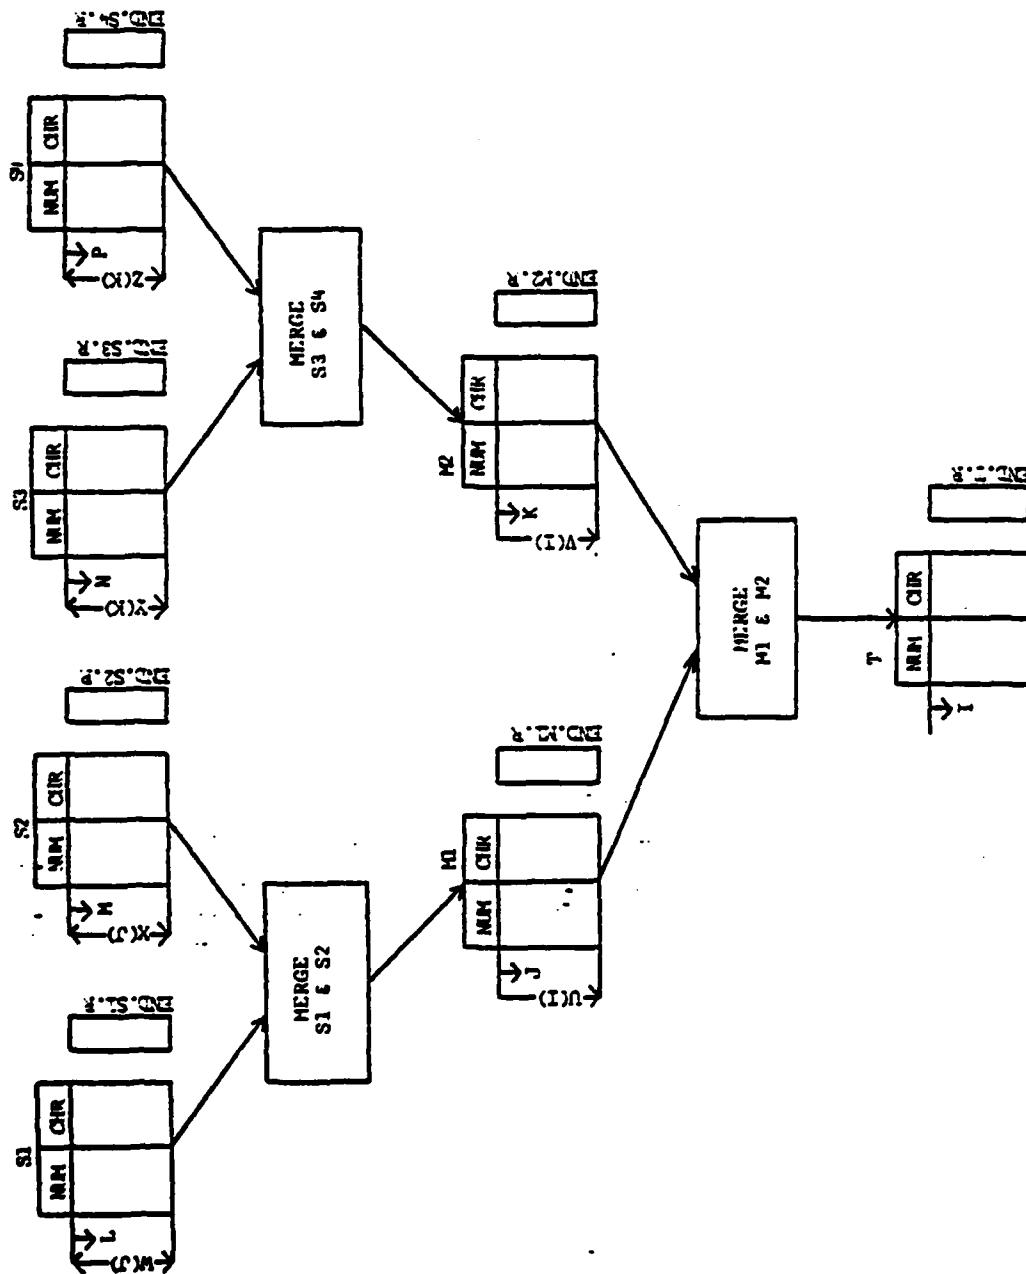
Fig. A.d Block Diagram of the Merging Example

The subscripts of the files in Fig. A.8 are shown as I, J, K, L, M, N, and P. The indirect subscripts for the latter six are U, V, W, X, Y, and Z, respectively. The definition of W(J) and X(J) is shown in the above specification in Fig. A.9 for the merger of S1 and S2.

```
/***** MERGE INPUT FILES S1 AND S2 INTO INTERIM FILE M1 *****/

END.S1.R(SUB1) = ENDFILE.S1.R(SUB1);

END.S2.R(SUB1) = ENDFILE.S2.R(SUB1);

XS1 IS GROUP (W(*));
   W IS FIELD (NUMERIC(4));

XS2 IS GROUP (X(*));
   X IS FIELD (NUMERIC(4));

DONES1 IS GROUP (DONES1F(*));
   DONES1F IS FIELD (BIT(1));

DONES2 IS GROUP (DONES2F(*));
   DONES2F IS FIELD (BIT(1));

SELS12 IS GROUP (SELS12F(*));
   SELS12F IS FIELD (BIT(1));

W(SUB1) = IF SUB1=1 THEN 1
          ELSE IF SELS12F(SUB1-1) & ^DONES1F(SUB1) THEN W(SUB1-1)+1
                                                    ELSE W(SUB1-1);

X(SUB1) = IF SUB1=1 THEN 1
          ELSE IF SELS12F(SUB1-1) ! DONES2F(SUB1) THEN X(SUB1-1)
                                                  ELSE X(SUB1-1)+1;

DONES1F(SUB1) = IF SUB1=1 THEN '0'B
                ELSE DONES1F(SUB1-1) !
                     (END.S1.R(W(SUB1-1)) & SELS12F(SUB1-1));

DONES2F(SUB1) = IF SUB1=1 THEN '0'B
                ELSE DONES2F(SUB1-1) !
                     (END.S2.R(X(SUB1-1)) & ^SELS12F(SUB1-1));

SELS12F(SUB1) = DONES2F(SUB1) !
                (^DONES1F(SUB1) & (S1.NUM(W(SUB1)) < S2.NUM(X(SUB1))));

M1.NUM(SUB1) = IF SELS12F(SUB1) THEN S1.NUM(W(SUB1))
                                ELSE S2.NUM(X(SUB1));

M1.CHR(SUB1) = IF SELS12F(SUB1) THEN S1.CHR(W(SUB1))
                                ELSE S2.CHR(X(SUB1));

END.M1.R(SUB1) = (DONES1F(SUB1) & END.S2.R(X(SUB1))) !
                 (DONES2F(SUB1) & END.S1.R(W(SUB1)));
*
```

Fig. A.9 MODEL specification for merging two files
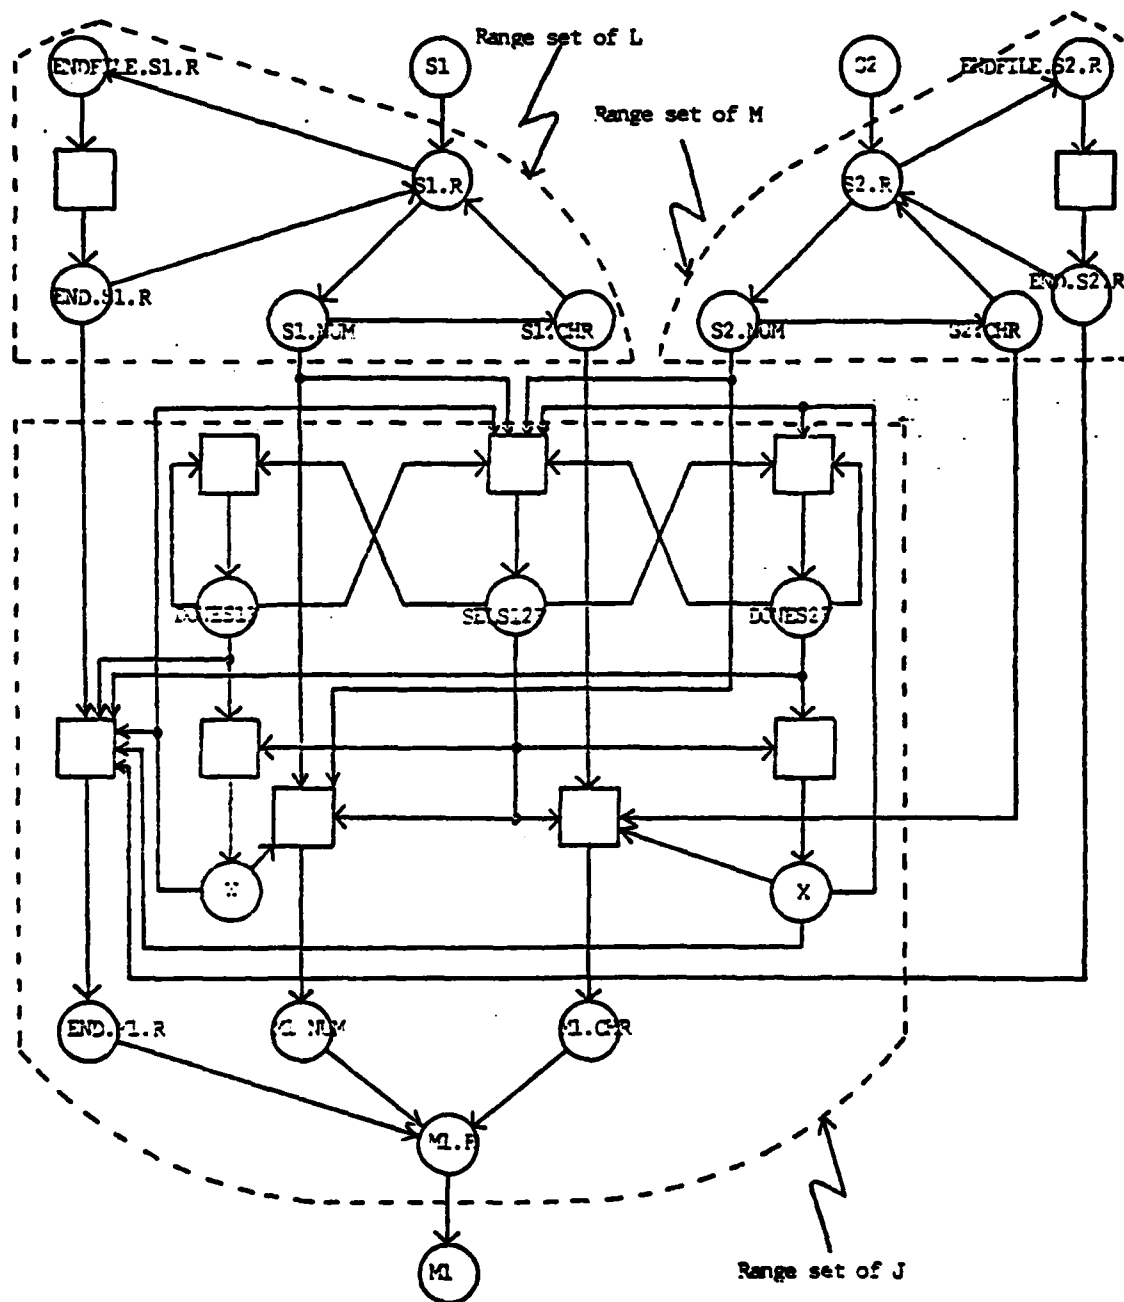
Fig. A.10 Array Graph for Merging two files

The entire specification is given in Fig. A.11.   There
are seven range sets that need to be merged progressively
(three for each of the mergers into M1 and M2, one for the
merger into T) into a single loop scope. The resulting
schedule is shown in Fig. A.12. The merger of range sets is
applied recursively resulting in nested conditional blocks
in the scope of the loop. Thus there are conditional blocks
for each of the source files of each merger, S1 and S2 into
M1 and S3 and S4 into M2. Further the conditional blocks of
these mergers are nested in the conditional blocks for
merging M1 and M2.   These conditional blocks are shown
bracketed in Fig. A.12.

```
/**********************************************************************/
/*                                                                    */
/* THE FOLLOWING SPECIFICATION DESCRIBES THE TARGET FILET1, WHICH IS  */
/* OBTAINED BY MERGING THE FOUR SORTED SOURCE FILES S1, S2, S3 AND S4.*/
/* THE MERGING IS DONE IN TWO STEPS.  FIRST, THE FILES S1 AND S2 ARE  */
/* MERGED INTO INTERIM FILE M1, AND THE FILES S3 AND S4 INTO M2.      */
/* M1 AND M2 ARE THEN MERGED INTO T.                                  */
/*                                                                    */
/**********************************************************************/


MODULE : MERGE4;
SOURCE : S1,S2,S3,S4;
TARGET : T;


S1 IS FILE (R(*));
   R IS RECORD (NUM,CHR);
      NUM IS FIELD (NUMERIC(4));
      CHR IS FIELD (CHAR(4));

S2 IS FILE (R(*)),

S3 IS FILE (R(*));

S4 IS FILE (R(*));

T IS FILE (R(*));

M1 IS FILE (R(*));

M2 IS FILE (R(*));


/***** SIZES OF INPUT FILES *****/

END.S1.R(SUB1) = ENDFILE.S1.R(SUB1);
END.S2.R(SUB1) = ENDFILE.S2.R(SUB1);
END.S3.R(SUB1) = ENDFILE.S3.R(SUB1);
END.S4.R(SUB1) = ENDFILE.S4.R(SUB1);
```

Fig. A.11 MODEL specification for merging four files

```
/***** MERGE INPUT FILES S1 AND S2 INTO INTERIM FILE M1 *****/

XS1 IS GROUP (W(*));
   W IS FIELD (NUMERIC(4));

XS2 IS GROUP (X(*));
   X IS FIELD (NUMERIC(4));

DONES1 IS GROUP (DONES1F(*));
   DONES1F IS FIELD (BIT(1));

DONES2 IS GROUP (DONES2F(*));
   DONES2F IS FIELD (BIT(1));

SELS12 IS GROUP (SELS12F(*));
   SELS12F IS FIELD (BIT(1));

W(SUB1) = IF SUB1=1 THEN 1
          ELSE IF SELS12F(SUB1-1) & ^DONES1F(SUB1) THEN W(SUB1-1)+1
                                                   ELSE W(SUB1-1);

X(SUB1) = IF SUB1=1 THEN 1
          ELSE IF SELS12F(SUB1-1) | DONES2F(SUB1) THEN X(SUB1-1)
                                                  ELSE X(SUB1-1)+1;

DONES1F(SUB1) = IF SUB1=1 THEN '0'B
                ELSE DONES1F(SUB1-1) |
                     (END.S1.R(W(SUB1-1)) & SELS12F(SUB1-1));

DONES2F(SUB1) = IF SUB1=1 THEN '0'B
                ELSE DONES2F(SUB1-1) |
                     (END.S2.R(X(SUB1-1)) & ^SELS12F(SUB1-1));

SELS12F(SUB1) = DONES2F(SUB1) |
                (^DONES1F(SUB1) & (S1.NUM(W(SUB1)) < S2.NUM(X(SUB1))));

M1.NUM(SUB1) = IF SELS12F(SUB1) THEN S1.NUM(W(SUB1))
                                ELSE S2.NUM(X(SUB1));

M1.CHR(SUB1) = IF SELS12F(SUB1) THEN S1.CHR(W(SUB1))
                                ELSE S2.CHR(X(SUB1));

END.M1.R(SUB1) = (DONES1F(SUB1) & END.S2.R(X(SUB1))) |
                 (DONES2F(SUB1) & END.S1.R(W(SUB1)));
```

Fig. A.11 MODEL specification for merging
four files(continued)

```
/***** MERGE INPUT FILES S3 AND S4 INTO INTERIM FILE M2 *****/

XS3 IS GROUP (Y(*));
   Y IS FIELD (NUMERIC(4));

XS4 IS GROUP (Z(*));
   Z IS FIELD (NUMERIC(4));

DONES3 IS GROUP (DONES3F(*));
   DONES3F IS FIELD (BIT(1));

DONES4 IS GROUP (DONES4F(*));
   DONES4F IS FIELD (BIT(1));

 ELS34 IS GROUP (SELS34F(*));
   SELS34F IS FIELD (BIT(1));

Y(SUB1) = IF SUB1=1 THEN 1
           ELSE IF SELS34F(SUB1-1) & ^DONES3F(SUB1) THEN Y(SUB1-1)+1
                                                     ELSE Y(SUB1-1);

Z(SUB1) = IF SUB1=1 THEN 1
           ELSE IF SELS34F(SUB1-1) | DONES4F(SUB1) THEN Z(SUB1-1)
                                                    ELSE Z(SUB1-1)+1;

DONES3F(SUB1) = IF SUB1=1 THEN '0'B
                 ELSE DONES3F(SUB1-1) |
                      (END.S3.R(Y(SUB1-1)) & SELS34F(SUB1-1));

DONES4F(SUB1) = IF SUB1=1 THEN '0'B
                 ELSE DONES4F(SUB1-1) |
                      (END.S4.R(Z(SUB1-1)) & ^SELS34F(SUB1-1));

SELS34F(SUB1) = DONES4F(SUB1) |
                 (^DONES3F(SUB1) & (S3.NUM(Y(SUB1)) < S4.NUM(Z(SU31))));

M2.NUM(SUB1) = IF SELS34F(SUB1) THEN S3.NUM(Y(SUB1))
                                ELSE S4.NUM(Z(SUB1));

M2.CHR(SUB1) = IF SELS34F(SUB1) THEN S3.CHR(Y(SUB1))
                                ELSE S4.CHR(Z(SUB1));

END.M2.R(SUB1) = (DONES3F(SUB1) & END.S4.R(Z(SUB1))) |
                 (DONES4F(SUB1) & END.S3.R(Y(SUB1)));
```

**Fig. A.11 MODEL specification for merging
four files(continued)**

```
/***** MERGE INTERIM FILES M1 AND M2 INTO OUTPUT FILE T1 *****/

XM1 IS GROUP (U(*));
   U IS FIELD (NUMERIC(4));

XM2 IS GROUP (V(*));
   V IS FIELD (NUMERIC(4));

DONEM1 IS GROUP (DONEM1F(*));
   DONEM1F IS FIELD (BIT(1));

DONEM2 IS GROUP (DONEM2F(*));
   DONEM2F IS FIELD (BIT(1));

SELM12 IS GROUP (SELM12F(*));
   SELM12F IS FIELD (BIT(1));

U(SUB1) = IF SUB1=1 THEN 1
          ELSE IF SELM12F(SUB1-1) & ^DONEM1F(SUB1) THEN U(SUB1-1)+1
                                                    ELSE U(SUB1-1);

V(SUB1) = IF SUB1=1 THEN 1
          ELSE IF SELM12F(SUB1-1) | DONEM2F(SUB1) THEN V(SUB1-1)
                                                  ELSE V(SUB1-1)+1;

DONEM1F(SUB1) = IF SUB1=1 THEN '0'B
                ELSE DONEM1F(SUB1-1) |
                    (END.M1.R(U(SUB1-1)) & SELM12F(SUB1-1));

DONEM2F(SUB1) = IF SUB1=1 THEN '0'B
                ELSE DONEM2F(SUB1-1) |
                    (END.M2.R(V(SUB1-1)) & ^SELM12F(SUB1-1));

SELM12F(SUB1) = DONEM2F(SUB1) |
                (^DONEM1F(SUB1) & (M1.NUM(U(SUB1)) < M2.NUM(V(SUB1))));

T.NUM(SUB1) = IF SELM12F(SUB1) THEN M1.NUM(U(SUB1))
                               ELSE M2.NUM(V(SUB1));

T.CHR(SUB1) = IF SELM12F(SUB1) THEN M1.CHR(U(SUB1))
                               ELSE M2.CHR(V(SUB1));

END.T.R(SUB1) = (DONEM1F(SUB1) & END.M2.R(V(SUB1))) |
                (DONEM2F(SUB1) & END.M1.R(U(SUB1)));
```

Fig. A.11 MODEL specification for merging
four files(continued)

Fig. A.12 Schedule for the MERGE4 Example

Fig. A.12 Schedule for the MERGE4 Example(continued)

```
MERGE4: PROCEDURE OPTIONS(MAIN);
DCL S1S RECORD SEQL INPUT;
DCL $FSTS1S BIT(1) INIT('1'B);
DCL ENDFILE$S1S BIT(1) INIT('0'B);
DCL S1_S CHAR(8) VARYING INIT('');
DCL S1_INDX FIXED BIN;
DCL S2S RECORD SEQL INPUT;
DCL $FSTS2S BIT(1) INIT('1'B);
DCL ENDFILE$S2S BIT(1) INIT('0'B);
DCL S2_S CHAR(8) VARYING INIT('');
DCL S2_INDX FIXED BIN;
DCL S3S RECORD SEQL INPUT;
DCL $FSTS3S BIT(1) INIT('1'B);
DCL ENDFILE$S3S BIT(1) INIT('0'B);
DCL S3_S CHAR(8) VARYING INIT('');
DCL S3_INDX FIXED BIN;
DCL S4S RECORD SEQL INPUT;
DCL $FSTS4S BIT(1) INIT('1'B);
DCL ENDFILE$S4S BIT(1) INIT('0'B);
DCL S4_S CHAR(8) VARYING INIT('');
DCL S4_INDX FIXED BIN;
DCL ($X2,$R_INTERIM$U) FIXED BIN;
DCL $B_INTERIM$U BIT(1);
DCL ($X3,$R_INTERIM$X) FIXED BIN;
DCL $B_INTERIM$X BIT(1);
DCL S2_R_S CHAR(8) VARYING;
DCL S2_R_INDX FIXED BIN;
DCL ($X4,$R_INTERIM$W) FIXED BIN;
DCL $B_INTERIM$W BIT(1);
DCL S1_R_S CHAR(8) VARYING;
DCL S1_R_INDX FIXED BIN;
DCL ($X5,$R_INTERIM$V) FIXED BIN;
DCL $B_INTERIM$V BIT(1);
DCL ($X6,$R_INTERIM$Z) FIXED BIN;
DCL $B_INTERIM$Z BIT(1);
DCL S4_R_S CHAR(8) VARYING;
DCL S4_R_INDX FIXED BIN;
DCL ($X7,$R_INTERIM$Y) FIXED BIN;
DCL $B_INTERIM$Y BIT(1);
DCL S3_R_S CHAR(8) VARYING;
DCL S3_R_INDX FIXED BIN;
DCL T_R_S CHAR(8) VARYING;
DCL T_R_S_F CHAR(8);
DCL T_R_SC BIT(64) BASED(ADDR(T_R_S_F));
DCL T_R_INDX FIXED BIN;
DCL TT RECORD SEQL OUTPUT;
DCL $FSTTT BIT(1) INIT('1'B);
DCL $ERROR_BUF CHAR(270) VAR;
DCL ERRORF FILE RECORD OUTPUT;
DCL ERRORF_BIT BIT(1) STATIC INIT('1'B);
DCL ($ERROR,$ACC_ERROR,$NOT_DONE)(20) BIT(1);
DCL $ERR_LAB(20) LABEL;
DCL $ERRSP$ FIXED BIN STATIC INITIAL (0);
DCL $TMP_VAL FLOAT BIN;
DCL $TMP_EPR BIT(1);
DECLARE
   1 M1,
     2 R,
       3 NUM(2) PIC'9999',
       3 CHR(2) CHAR(4);
```

Fig. A.13 Generated PL/I Program for
the MERGE4 Example

```
DECLARE
   1 M2,
    2 R,
     3 NUM(2) PIC'9999',
     3 CHR(2) CHAR(4);
DECLARE
   1 S1,
    2 R,
     3 NUM(2) PIC'9999',
     3 CHR(2) CHAR(4);
DECLARE
   1 S2,
    2 R,
     3 NUM(2) PIC'9999',
     3 CHR(2) CHAR(4);
DECLARE
   1 S3,
    2 R,
     3 NUM(2) PIC'9999',
     3 CHR(2) CHAR(4);
DECLARE
   1 S4,
    2 R,
     3 NUM(2) PIC'9999',
     3 CHR(2) CHAR(4);
DECLARE
   1 T,
    2 R,
     3 NUM PIC'9999',
     3 CHR CHAR(4);
DECLARE
   1 INTERIM,
    2 XS4,
     3 Z(2) PIC'9999',
    2 XS3,
     3 Y(2) PIC'9999',
    2 XS2,
     3 X(2) PIC'9999',
    2 XS1,
     3 W(2) PIC'9999',
    2 XM2,
     3 V(2) PIC'9999',
    2 XM1,
     3 U(2) PIC'9999',
    2 SELS34,
     3 SELS34F(2) BIT(1),
    2 SELS12,
     3 SELS12F(2) BIT(1),
    2 SELM12,
     3 SELM12F(2) BIT(1),
    2 DONES4,
     3 DONES4F(2) BIT(1),
    2 DONES3,
     3 DONES3F(2) BIT(1),
    2 DONES2,
     3 DONES2F(2) BIT(1),
    2 DONES1,
     3 DONES1F(2) BIT(1),
    2 DONEM2,
     3 DONEM2F(2) BIT(1),
```

Fig. A.13 Generated PL/I Program for
the MERGE4 Example(Continued)

```
      2 DONEM1,
        3 DONEM1F(2) BIT(1),
       2 $YSGEN1,
        3 ENDST_R(2) BIT(1) ,
       2 $YSGEN2,
        3 END$M2_R(2) BIT(1) ,
       2 $YSGEN3,
        3 END$M1_R(2) BIT(1) ,
       2 $YSGEN4,
        3 END$S4_R(2) BIT(1) ,
       2 $YSGEN5,
        3 END$S3_R(2) BIT(1) ,
       2 $YSGEN6,
        3 END$S2_R(2) BIT(1) ,
       2 $YSGEN7,
        3 END$S1_R(2) BIT(1) ,
       2 $YSGEN8,
        3 ENDFILE$S4_R BIT(1) ,
       2 $YSGEN9,
        3 ENDFILE$S3_R BIT(1) ,
       2 $YSGEN10,
        3 ENDFILE$S2_R BIT(1) ,
       2 $YSGEN11,
        3 ENDFILE$S1_R BIT(1) ;
DCL $I1 FIXED BIN;
DCL (TRUE,SELECTED) BIT(1) INIT('1'B);
DCL (FALSE,NOT_SELE,NOT_SELECTED) BIT(1) INIT('0'B);
ON ENDFILE($1S) BEGIN;
ENDFILE$S1S='1'B;
 S1_S=COPY(' ',8);
 END;
ON ENDFILE($2S) BEGIN;
ENDFILE$S2S='1'B;
 S2_S=COPY(' ',8);
 END;
ON ENDFILE($3S) BEGIN;
ENDFILE$S3S='1'B;
 S3_S=COPY(' ',8);
 END;
ON ENDFILE($4S) BEGIN;
ENDFILE$S4S='1'B;
 S4_S=COPY(' ',8);
 END;
ON UNDEFINEDFILE(ERRORF)  ERRORF_BIT='0'B;
/* ON ERROR BEGIN;
IF ERRORF_BIT THEN WRITE FILE(ERRORF) FROM ($ERROR_BUF);
$ERROR($ERRSP$)='1'B;
$ACC_ERROR($ERRSP$)='1'B;
GO TO $ERR_LAB($ERRSP$) ;
  END ;
  */
ERROR_RESTART: ;
$ERRSP$ = $ERRSP$ +1;
$ERROR($ERRSP$)='0'B;
$ACC_ERROR($ERRSP$)='0'B;
$ERR_LAB($ERRSP$)=END_PROGRAM;
OPEN FILE($1S);
OPEN FILE($2S);
OPEN FILE($3S);
OPEN FILE($4S);
```

Fig. A.13 Generated PL/I Program for
the MERGE4 Example(Continued)

```
$ERRSP$ = $ERRSP$ +1;
$ERROR($ERRSP$)='0'B;
$ACC_ERROR($ERRSP$)='0'B;
$ERR_LAB($ERRSP$)=LOOP_END1;
$I1 =0;
$NOT_DONE(1)='1'B;
DO  WHILE($NOT_DONE(1));
  $I1 = $I1 +1;
  $ERROR($ERRSP$)='0'B;
  IF $I1=1 THEN INTERIM.DONEM1F(2)='0'B;
   ELSE INTERIM.DONEM1F(2)=INTERIM.DONEM1F(1);END$M1_R(1-INTERIM.U(1)+INTERIM.U(
  1))&INTERIM.SELM12F(1);
  IF $I1=1 THEN INTERIM.U(2)=1;
   ELSE IF INTERIM.SELM12F(1)&^INTERIM.DONEM1F(2) THEN INTERIM.U(2)=INTERIM.U(1)
  +1;
   ELSE INTERIM.U(2)=INTERIM.U(1);
  IF $I1=1 THEN DO;
   $B_INTERIMSU= '1'B;
   $R_INTERIMSU=0;
   END;
   ELSE IF (INTERIM.U(2)>INTERIM.U(1)) THEN DO;
   $B_INTERIMSU= '1'B;
   $R_INTERIMSU=0;
   END;
   ELSE DO;
   $B_INTERIMSU= '0'B;
   $R_INTERIMSU=1;
   END;
  IF $I1=1 THEN INTERIM.DONEM2F(2)='0'B;
   ELSE INTERIM.DONEM2F(2)=INTERIM.DONEM2F(1);END$M2_R(1-INTERIM.V(1)+INTERIM.V(
  1))&^INTERIM.SELM12F(1);
  IF $I1=1 THEN INTERIM.V(2)=1;
   ELSE IF INTERIM.SELM12F(1);INTERIM.DONEM2F(2) THEN INTERIM.V(2)=INTERIM.V(1);
   ELSE INTERIM.V(2)=INTERIM.V(1)+1;
  IF $I1=1 THEN DO;
   $B_INTERIMSV= '1'B;
   $R_INTERIMSV=0;
   END;
   ELSE IF (INTERIM.V(2)>INTERIM.V(1)) THEN DO;
   $B_INTERIMSV= '1'B;
   $R_INTERIMSV=0;
   END;
   ELSE DO;
   $B_INTERIMSV= '0'B;
   $R_INTERIMSV=1;
   END;
  IF $B_INTERIMSU THEN
  DO ;
    $X2 = INTERIM.U(2);
    IF $X2=1 THEN INTERIM.DONES1F(2)='0'B;
     ELSE INTERIM.DONES1F(2)=INTERIM.DONES1F(1);END$S1_R(1-INTERIM.W(1)
    +INTERIM.W(1))&INTERIM.SELS12F(1);
    IF $X2=1 THEN INTERIM.W(2)=1;
     ELSE IF INTERIM.SELS12F(1)&^INTERIM.DONES1F(2) THEN INTERIM.W(2)=INTERIM.W(
    1)+1;
     ELSE INTERIM.W(2)=INTERIM.W(1);
    IF $X2=1 THEN DO;
     $B_INTERIMSW= '1'B;
     $R_INTERIMSW=0;
     END;
```

Fig. A.13 Generated PL/I Program for
the MERGE4 Example(Continued)

```
    ELSE IF (INTERIM.W(2)>INTERIM.W(1)) THEN DO;
    $B_INTERIMSW= '1'B;
    $R_INTERIMSW=0;
    END;
    ELSE DO;
    $B_INTERIMSW= '0'B;
    $R_INTERIMSW=1;
    END;
IF $X2=1 THEN INTERIM.DONES2F(2)='0'B;
    ELSE INTERIM.DONES2F(2)=INTERIM.DONES2F(1)!END$S2_R(1-INTERIM.X(1)
+INTERIM.X(1))&^INTERIM.SELS12F(1);
IF $X2=1 THEN INTERIM.X(2)=1;
    ELSE IF INTERIM.SELS12F(1)!INTERIM.DONES2F(2) THEN INTERIM.X(2)=INTERIM.X(
1);
    ELSE INTERIM.X(2)=INTERIM.X(1)+1;
IF $X2=1 THEN DO;
    $B_INTERIMSX= '1'B;
    $R_INTERIMSX=0;
    END;
    ELSE IF (INTERIM.X(2)>INTERIM.X(1)) THEN DO;
    $B_INTERIMSX= '1'B;
    $R_INTERIMSX=0;
    END;
    ELSE DO;
    $B_INTERIMSX= '0'B;
    $R_INTERIMSX=1;
    END;
IF $B_INTERIMSX THEN
DO ;
  $X3 = INTERIM.X(2);
  IF $FSTS2S THEN DO;
    READ FILE(S2S) INTO (S2_R_S);
    $FSTS2S='0'B;
    END;
  ELSE S2_R_S=S2_S;
  S2_R_INDX=1;
  IF ^ENDFILE$S2S THEN READ FILE(S2S) INTO (S2_S);
  $ERROR_BUF=S2_R_S;
  ENDFILE$S2_R=ENDFILE$S2S;
  END$S2_R(2)=ENDFILE$S2_R;
  UNSPEC(S2.NUM(2))=UNSPEC(SUBSTR(S2_R_S,S2_R_INDX,4)) ;
  S2_R_INDX=S2_R_INDX+4 ;
  S2.CHR(2)=SUBSTR(S2_R_S,S2_R_INDX,4) ;
  S2_R_INDX=S2_R_INDX+4 ;
END;
IF $B_INTERIMSW THEN
DO ;
  $X4 = INTERIM.W(2);
  IF $FSTS1S THEN DO;
    READ FILE(S1S) INTO (S1_R_S);
    $FSTS1S='0'B;
    END;
  ELSE S1_R_S=S1_S;
  S1_R_INDX=1;
  IF ^ENDFILE$S1S THEN READ FILE(S1S) INTO (S1_S);
  $ERROR_BUF=S1_R_S;
  ENDFILE$S1_R=ENDFILE$S1S;
  END$S1_R(2)=ENDFILE$S1_R;
  UNSPEC(S1.NUM(2))=UNSPEC(SUBSTR(S1_R_S,S1_R_INDX,4)) ;
  S1_R_INDX=S1_R_INDX+4 ;
```

Fig. A.13 Generated PL/I Program for
the MERGE4 Example(Continued)

```
      S1.CHR(2)=SUBSTR(S1_R_S,S1_R_INDX,4) ;
      S1_R_INDX=S1_R_INDX+4 ;
    END;
    INTERIM.SELS12F(2)=INTERIM.DONES2F(2)!^INTERIM.DONES1F(2)&S1.NUM(
    2-$R_INTERIM$W)<S2.NUM(2-$R_INTERIM$X);
    IF INTERIM.SELS12F(2) THEN M1.CHR(2)=S1.CHR(2-$R_INTERIM$W);
     ELSE M1.CHR(2)=S2.CHR(2-$R_INTERIM$X);
    IF INTERIM.SELS12F(2) THEN M1.NUM(2)=S1.NUM(2-$R_INTERIM$W);
     ELSE M1.NUM(2)=S2.NUM(2-$R_INTERIM$X);
    END$M1_R(2)=INTERIM.DONES1F(2)&END$S2_R(2-$R_INTERIM$X)!INTERIM.DONES2F(2)
    &END$S1_R(2-$R_INTERIM$W);
    IF $B_INTERIM$X THEN S2.CHR(1) = S2.CHR(2);
    IF $B_INTERIM$X THEN S2.NUM(1) = S2.NUM(2);
    IF $B_INTERIM$X THEN END$S2_R(1) = END$S2_R(2);
    IF $B_INTERIM$W THEN S1.CHR(1) = S1.CHR(2);
    IF $B_INTERIM$W THEN S1.NUM(1) = S1.NUM(2);
    IF $B_INTERIM$W THEN END$S1_R(1) = END$S1_R(2);
  END;
  IF $B_INTERIM$V THEN
  DO ;
    $X5 = INTERIM.V(2);
    IF $X5=1 THEN INTERIM.DONES3F(2)='0'B;
     ELSE INTERIM.DONES3F(2)=INTERIM.DONES3F(1)!END$S3_R(1-INTERIM.Y(1)
    +INTERIM.Y(1))&INTERIM.SELS34F(1);
    IF $X5=1 THEN INTERIM.Y(2)=1;
     ELSE IF INTERIM.SELS34F(1)&^INTERIM.DONES3F(2) THEN INTERIM.Y(2)=INTERIM.Y(
    1)+1;
     ELSE INTERIM.Y(2)=INTERIM.Y(1);
    IF $X5=1 THEN DO;
     $B_INTERIM$Y= '1'B;
     $R_INTERIM$Y=0;
     END;
     ELSE IF (INTERIM.Y(2)>INTERIM.Y(1)) THEN DO;
     $B_INTERIM$Y= '1'B;
     $R_INTERIM$Y=0;
     END;
     ELSE DO;
     $B_INTERIM$Y= '0'B;
     $R_INTERIM$Y=1;
     END;
    IF $X5=1 THEN INTERIM.DONES4F(2)='0'B;
     ELSE INTERIM.DONES4F(2)=INTERIM.DONES4F(1)!END$S4_R(1-INTERIM.Z(1)
    +INTERIM.Z(1))&^INTERIM.SELS34F(1);
    IF $X5=1 THEN INTERIM.Z(2)=1;
     ELSE IF INTERIM.SELS34F(1)!INTERIM.DONES4F(2) THEN INTERIM.Z(2)=INTERIM.Z(
    1);
     ELSE INTERIM.Z(2)=INTERIM.Z(1)+1;
    IF $X5=1 THEN DO;
     $B_INTERIM$Z= '1'B;
     $R_INTERIM$Z=0;
     END;
     ELSE IF (INTERIM.Z(2)>INTERIM.Z(1)) THEN DO;
     $B_INTERIM$Z= '1'B;
     $R_INTERIM$Z=0;
     END;
     ELSE DO;
     $B_INTERIM$Z= '0'B;
     $R_INTERIM$Z=1;
     END;
    IF $B_INTERIM$Z THEN
```

Fig. A.13 Generated PL/I Program for
the MERGE4 Example(Continued)

```
DO ;
  $X6 = INTERIM.Z(2);
  IF $FSTS4S THEN DO;
   READ FILE(S4S) INTO (S4_R_S);
   $FSTS4S='0'B;
   END;
  ELSE S4_R_S=S4_S;
  S4_R_INDX=1;
  IF ^ENDFILE$S4S THEN READ FILE(S4S) INTO (S4_S);
  $ERROR_BUF=S4_R_S;
  ENDFILE$S4_R=ENDFILE$S4S;
  END$S4_R(2)=ENDFILE$S4_R;
  UNSPEC(S4.NUM(2))=UNSPEC(SUBSTR(S4_R_S,S4_R_INDX,4)) ;
  S4_R_INDX=S4_R_INDX+4 ;
  S4.CHR(2)=SUBSTR(S4_R_S,S4_R_INDX,4) ;
  S4_R_INDX=S4_R_INDX+4 ;
END;
IF $B_INTERIM$Y THEN
DO ;
  $X7 = INTERIM.Y(2);
  IF $FSTS3S THEN DO;
   READ FILE(S3S) INTO (S3_R_S);
   $FSTS3S='0'B;
   END;
  ELSE S3_R_S=S3_S;
  S3_R_INDX=1;
  IF ^ENDFILE$S3S THEN READ FILE(S3S) INTO (S3_S);
  $ERROR_BUF=S3_R_S;
  ENDFILE$S3_R=ENDFILE$S3S;
  END$S3_R(2)=ENDFILE$S3_R;
  UNSPEC(S3.NUM(2))=UNSPEC(SUBSTR(S3_R_S,S3_R_INDX,4)) ;
  S3_R_INDX=S3_R_INDX+4 ;
  S3.CHR(2)=SUBSTR(S3_R_S,S3_R_INDX,4) ;
  S3_R_INDX=S3_R_INDX+4 ;
END;
INTERIM.SELS34F(2)=INTERIM.DONES4F(2)!^INTERIM.DONES3F(2)&S3.NUM(
2-$R_INTERIM$Y)<S4.NUM(2-$R_INTERIM$Z);
IF INTERIM.SELS34F(2) THEN M2.CHR(2)=S3.CHR(2-$R_INTERIM$Y);
 ELSE M2.CHR(2)=S4.CHR(2-$R_INTERIM$Z);
IF INTERIM.SELS34F(2) THEN M2.NUM(2)=S3.NUM(2-$R_INTERIM$Y);
 ELSE M2.NUM(2)=S4.NUM(2-$R_INTERIM$Z);
END$M2_R(2)=INTERIM.DONES3F(2)&END$S4_R(2-$R_INTERIM$Z)!INTERIM.DONES4F(2)
&END$S3_R(2-$R_INTERIM$Y);
IF $B_INTERIM$Z THEN S4.CHR(1) = S4.CHR(2);
IF $B_INTERIM$Z THEN S4.NUM(1) = S4.NUM(2);
IF $B_INTERIM$Z THEN END$S4_R(1) = END$S4_R(2);
IF $B_INTERIM$Y THEN S3.CHR(1) = S3.CHR(2);
IF $B_INTERIM$Y THEN S3.NUM(1) = S3.NUM(2);
IF $B_INTERIM$Y THEN END$S3_R(1) = END$S3_R(2);
END;
END$T_F(2)=INTERIM.DONEM1F(2)&END$M2_R(2-$R_INTERIM$V)!INTERIM.DONEM2F(2)
&END$M1_R(2-$R_INTERIM$U);
INTERIM.SELM12F(2)=INTERIM.DONEM2F(2)!^INTERIM.DONEM1F(2)&M1.NUM(
2-$R_INTERIM$U)<M2.NUM(2-$R_INTERIM$V);
IF INTERIM.SELM12F(2) THEN T.CHR=M1.CHR(2-$R_INTERIM$U);
 ELSE T.CHR=M2.CHR(2-$R_INTERIM$V);
IF INTERIM.SELM12F(2) THEN T.NUM=M1.NUM(2-$R_INTERIM$U);
 ELSE T.NUM=M2.NUM(2-$R_INTERIM$V);
T_R_INDX=1;
SUBSTR(T_R_SC,T_R_INDX*8-7,4*8)=UNSPEC(T.NUM) ;
```

Fig. A.13 Generated PL/I Program for
the MERGE4 Example(Continued)

```
T_R_INDX=T_R_INDX+4 ;
SUBSTR(T_R_S_F,T_R_INDX,4)=T.CHR ;
T_R_INDX=T_R_INDX+4 ;
T_R_S=SUBSTR(T_R_S_F,1,T_R_INDX-1);
WRITE FILE(TT) FROM (T_R_S);
LOOP_END1;;
IF ENDST_R(2) THEN $NOT_DONE(1)='0'B;
INTERIM.SELM12F(1) = INTERIM.SELM12F(2);
ENDST_R(1) = ENDST_R(2);
INTERIM.V(1) = INTERIM.V(2);
INTERIM.DONEM2F(1) = INTERIM.DONEM2F(2);
INTERIM.U(1) = INTERIM.U(2);
INTERIM.DONEM1F(1) = INTERIM.DONEM1F(2);
IF $B_INTERIMSU THEN ENDSM1_R(1) = ENDSM1_R(2);
IF $B_INTERIMSU THEN M1.NUM(1) = M1.NUM(2);
IF $B_INTERIMSU THEN M1.CHR(1) = M1.CHR(2);
IF $B_INTERIMSU THEN INTERIM.SELS12F(1) = INTERIM.SELS12F(2);
IF $B_INTERIMSU THEN INTERIM.X(1) = INTERIM.X(2);
IF $B_INTERIMSU THEN INTERIM.DONES2F(1) = INTERIM.DONES2F(2);
IF $B_INTERIMSU THEN INTERIM.W(1) = INTERIM.W(2);
IF $B_INTERIMSU THEN INTERIM.DONES1F(1) = INTERIM.DONES1F(2);
IF $B_INTERIMSV THEN ENDSM2_R(1) = ENDSM2_R(2);
IF $B_INTERIMSV THEN M2.NUM(1) = M2.NUM(2);
IF $B_INTERIMSV THEN M2.CHR(1) = M2.CHR(2);
IF $B_INTERIMSV THEN INTERIM.SELS34F(1) = INTERIM.SELS34F(2);
IF $B_INTERIMSV THEN INTERIM.Z(1) = INTERIM.Z(2);
IF $B_INTERIMSV THEN INTERIM.DONES4F(1) = INTERIM.DONES4F(2);
IF $B_INTERIMSV THEN INTERIM.Y(1) = INTERIM.Y(2);
IF $B_INTERIMSV THEN INTERIM.DONES3F(1) = INTERIM.DONES3F(2);
END;
$TMP_ERR=$ACC_ERROR($ERRSP$);
$ERRSP$ = $ERRSP$ - 1;
IF $TMP_ERR THEN $ERROR($ERRSP$)='1'B;
IF $TMP_ERR THEN $ACC_ERROR($ERRSP$)='1'B;
CLOSE FILE(TT);
END_PROGRAM: RETURN;
END MERGE4;
$
```

Fig. A.13 Generated PL/I Program for
the MERGE4 Example(Continued)

INDEX

BIBLIOGRAPHY

ADAA 79

Preliminary ADA Reference Manual.
SIGPLAN Notices 14(6): Part A, June, 1979.

ADAB 79

Rationale for the Design of the ADA Programming
    Language.
SIGPLAN Notices 14(6): Part B, June, 1979.

AHJO 76

Aho, A.V., and Johnson, S.C.
Optimal Code Generation for Expression Trees.
JACM 23(3):488-501, July, 1976.

AHJO 77

Aho, A.V., and Johnson, S.C.
Code Generation for Expressions with Common
    Subexpressions.
JACM 24(1):146-160, January, 1977.

AHU 74

Aho, A.V., Hopcroft, J. E., and Ullman, J. D.
The Design and Analysis of Computer Algorithms.
Reading, Mass., Addison-Wesley, 1974.

AHUL 78

Aho, A.V., and Ullman, J.D.
Principles of Compiler Design.
Addison-Wesley, 1978.

ASWA 77

Ashcroft, E.A., and Wadge, W.W.
Lucid, A Non-procedural Language with Iteration.
CACM 20(7):519-526, July, 1977.

BAFI 79

Bauer, J., and Finger, A.
Test Plan Generation Using Formal Grammars.
Fourth International Conference on Software
    Engineering, pp. 425-432. September, 1979.

BAKA 76

Barstow, D.R., and Kant, E.
Observations on the Interaction Between Coding and
    Efficiency Knowledge in the PSI Program
    Synthesis System.
Second International Conference on Software
    Engineering, pp. 19-31. October, 1976.

BOKP 76

Bobillier, P. A., Kahan, B. C., and Probst, A. R.

Simulation with GPSS and GPSS V.
Prentice-Hall, 1976.

BRSE 76

Bruno, J., and Sethi, R.
Code Generation for a One-Register Machine.
JACM 23(3):502-510, July, 1976.

BUDA 77

Burstall, R. M., and Darlington, J.
A Transformation System for Developing Recursive
    Programs.
JACM 24(1):44-67, January, 1977.

CHTH 79

Cheatham, T. E., Townley, J. A., and Holloway, G. H.
A System for Program Refinement.
Fourth International Conference on Software
    Engineering, pp. 53-62. September, 1979.

CHHT 81

Cheatham, T. E., Holloway, G. H., and Townle. . J. A.
Program Refinement by Transformation.
Fifth International Conference on Software
    Engineering, pp. 430-437. March, 1981.

DAMN 70

Dahl, O. J., Myhrhaug, B., and Nygaard, K.
The SIMULA 67 Common Base Language.
Publication S-22, Norwegian Computing Center, Oslo,
1970.

DAVI 79

Davis, A., et al.
RLP: An Automated Tool for the Automatic Processing
    of Requirements.
Proceedings of COMPSAC 1979, pp. 289-299.

DAVI 80

Davis, A.
Automating the Requirements Phase: Benefits to
    Later Phases of the Software Life-Cycle.
Proceedings of COMPSAC 1980, pp. 42-48.?

FELD 72

Feldman, J.A.
Automatic Programming.
Technical Report STAN-CS-72-255, Stanford
University,
    February, 1972.

GEMS 77

Geschke, C.M., Morris, J.H., and Satterthwaite, E.H.

Early Experience with MESA.
CACM 20(8):540-553, August, 1977.

GOKH 81

Gokhale, M.
Data Flow Analysis of Non-procedural Specifications.
Proposal for PhD research. University of
    Pennsylvania, July, 1981.

GRAH 80

Graham, S. L.
Table-Driven Code Generation.
IEEE Computer, 13(8):25-34, August, 1980.

GREB 81

Greenburg, R.
Simultaneous Equations in the MODEL system with an
    Application to Econometric Modelling
Master Thesis, University of Pennsylvania, October,
1981.

GREE 69

Green, C.
Application of Theorem Proving to Problem Solving.
First Joint Conference on Artificial Intelligence,
    pp. 219-239. 1969.

GREE 77

Green, C.
The Design of the PSI Program Synthesis System.
Second International Conference on Software
    Engineering, pp. 4-18. IEEE Press, October, 1976.

HHKW 77

Hammer, M., Howe, W.G., Kruskal, V.J.,
    and Wladawsky, I.
A Very High Level Programming Language for Data
    Processing Applications.
CACM 20(11):832-840, November, 1977.

IVER 62

Iverson, K.E.
A Programming Language.
Wiley, New York, 1962.

KESC 75

Kennedy, K., and Schwartz, J.
An Introduction to the Set Theoretical Language
    SETL.
Comp. and Maths with Appls, Vol. 1, pp. 97-119.
    Pergamon Press, 1975. Great Britain.

KKPL 81

Kuck, D. J., Kuhn, R. H., Padua, D. A.,
    and Wolfe, M.
Dependence Graphs and Compiler Optimizations
In Proceedings of the ACM National Conference,
    Pages 207-218. ACM, 1981.

LCHN 80

Leverett, B. W., Cattell, R. G., Hobbs, S. O.,
    Newcomer, J. M., Reiner, A. H., Schatz, B. R.,
    and Wulf, W. A.
An Overview of the Production-Quality-Compiler-
    Compiler Project.
IEEE Computer, 13(8):38-49, August, 1980.

LEWA 74

Lee, R.C.T., and Waldinger, R.J.
An Improved Program Synthesis Algorithm and its
    Correctness.
CACM 17(4):211-217, April, 1974.

LSAS 77

Liskov, B., Snyder, A., Atkinson, R.,
    and Shaffert, C.
Abstraction Mechanisms in CLU.
CACM 20(8):564-574, August, 1977.

LYNC 69

Lynch, H.
ADS: A Technique in System Documentation.
Database, 1, 1 (Spring 1969), pp. 6-18.

MANN 71

Manna, Z.
Toward Automatic Program Synthesis.
CACM 14(3): 151-165, March, 1971.

MCCU 77

McCune, B.P.
The PSI Program Model Builder: Synthesis of Very
    High-Level Programs.
SIGPLAN Notices 12(8):130-138, August, 1977.

NUKO 76

Nunamaker, J.F., and Konsynski, B.R.
Computer-Aided Analysis and Design of Information
    System.
CACM 19(12):674-687, December, 1976.

NUNA 71

Nunamaker, J.F.
A Methodology for the Design and Optimization of
    Information Processing Systems.
SJCC, pp. 283-294, AFIPS Press, 1971.

MICROCOPY RESOLUTION TEST CHART
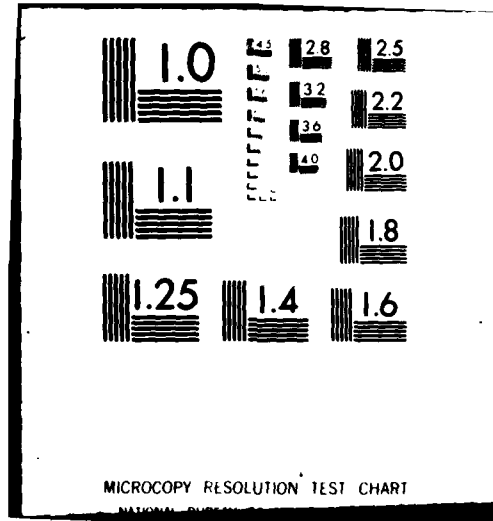
PARN 72

    Parnas, D.L.
    On the Criteria to be Used in Decomposing Systems
       into MOdules.
    CACM 15(12):1053-1058, December, 1972.

PHIL 77

    Phillips, J.V.
    Program Inference from Traces Using Multiple
       Knowledge Sources.
    Fifth Joint Conference on Artificial Intelligence,
       pp. 812. August, 1977.

PNPR 80

    Pnueli, A, and Prywes, N.
    Operations on Arrays and Data Structures.
    Technical Report, MCS-79-0298, University of
       Pennsylvania, 1980.

PNPR 81

    Pnueli, A, and Prywes, N.
    Distributed Processing in the MODEL System with an
       Application to Econometric Modelling.
    Technical Report, MCS-79-0298, University of
       Pennsylvania, 1981.

RAMI 73

    Ramirez, J.A.
    Automatic Generation of Data Conversion Programs
       Using A Data Definition Language.
    PhD thesis, University of Pennsylvania, 1973.

RIN 76

    Rin, N.A.
    Automatic Generation of Data Processing Programs
       from a Non-procedural Language.
    PhD thesis, University of Pennsylvania, 1976.

SANG 80

    Sangal, R.
    Modularity in Non-procedural Languages Through
       Abstract Data Types.
    PhD thesis, University of Pennsylvania, July, 1980.

SCAN 72

    Schneck, P. B., and Angel, E.
    FORTRAN to FORTRAN Optimizing Compiler.
    The Computer Journal (British Computer Society)
       16(4):322-329, 1972.

SCH 75

    Schwartz, J. T.

Optimization of Very High Level Languages-I.
Journal of Computer Languages, Vol. 1, pp. 161-194.
    Pergamon Press, 1975. Northern Ireland.

SCHA 73

Schaefer, M.
A Mathematical Theory of Global Program
    Optimization.
Prentice-Hall, New Jersey, 1973.

SCHW 75

Schwartz, J. T.
Optimization of Very High Level Languages-II.
Journal of Computer Languages, Vol. 1, pp. 197-218.
    Pergamon Press, 1975. Northern Ireland.

SHAS 78

Shastry, S.
Verification and Correction of Non-procedural
    Specification in Automatic Generation of
    Programs.
PhD thesis, University of Pennsylvania, 1978.

SZHM 77

Szolovits, P., Hawkinson, L., and Martin, W.
An Overview of OWL, A Language for Knowledge
    Representation.
Technical Report MIT-LCS-TM-86, MIT, June, 1977.

WALE 69

Waldinger, R.J., and Lee, R.C.T.
PROW: A Step Towards Automatic Program Writing.
First Joint Conference on Artificial Intelligence,
    pp. 241-252. 1969.

WEGN 79

Wegner, P.
Research Directions in Software Technology.
MIT Press, 1979.